THE COMPLETE
TECHNICAL PAPER PROCEEDINGS
FROM:

NCTA
Technical
Papers

A SIMPLE APPROACH FOR DERIVING THE SYMBOL ERROR RATE OF NON-RECTANGULAR 2^{2K+1} *M*-ARY AMPM MODULATION

Patricio Latini - Ayham Al-Banna

ARRIS Group

Abstract

This paper proposes a simple method to derive a closed-form expression for the exact symbol error probability of Non-Rectangular 2^{2k+1} Amplitude Modulated Phase Modulated (AMPM) signals over an Additive White Gaussian Noise (AWGN) channel. The obtained expression is verified using MATLAB-based computer simulations of AMPM systems for different QAM modulation orders, including QAM8, QAM32, and QAM128. Finally, the derived equation for the exact symbol error probability is compared with the upper-bounded symbol error probability expression for different QAM orders

.

I. INTRODUCTION

One way to achieve high data rates over band-limited channels is to increase the number of bits per symbol using optimal signal constellations designed to provide efficient performance. Nowadays, Quadrature Amplitude Modulation (QAM) is one of the most common modulation schemes used in communications systems. In particular, square QAM constellations [1] that contain even number of bits, *2k*, are widely used in many applications. Such systems have M symbols ($M=2^{2k}$), where the symbols are arranged to produce a square signal constellation.

Recently, the desire to reduce transmission errors in communications systems motivated burst receivers to support techniques like Trellis Coded Modulations (TCM) [1] in order to gain higher decoding granularity and therefore recover symbols more accurately. Since TCM mainly uses double number of symbols when compared to a standard QAM signal constellation that does not employ TCM, the need for QAM constellations with odd number of bits ($M=2^{2k+1}$) has elevated lately.

Odd-bit QAM constellations are currently used in many applications like DOCSIS and HDSL [2] [3] [4]. However, some of these applications use odd-bit QAM constellations that are not arranged in rectangular fashion. In particular, these QAM constellations are arranged such that they are a special case of Amplitude-Modulated Phase-Modulated (AMPM) signal constellations, which are designed to provide better efficiency in nonlinear distortion communication channels [5] [6]. AMPM modulation is also referred to as Carrierless Amplitude And Phase (CAP)-QAM modulation [4] [7].

Most of the previous research work has focused on rectangular even-bit QAM constellations, where a closed-form expression for the exact probability of symbol error in the presence of AWGN has already been analyzed and established [1] [8] [9] [10]. Researchers also studied odd-bit rectangular QAM constellations and developed expressions for the exact probability of symbol error after it was upper-bounded by the probability of symbol error of even-bit rectangular QAM constellations. In fact, various ways were developed, including simple geometrical procedures, to obtain a closed-form expression for the exact probability of symbol error in the presence of AWGN. Specifically, the authors in [11] proposed a geometrical approach to derive a closed-form expression for the probability of symbol error for the special case 8-symbol rectangular QAM system in the presence of AWGN. While simple geometrical approaches were used to analyze even-bit and some odd-bit rectangular QAM systems, no equivalent work has been done for the more complicated odd-bit AMPM systems, where the signal constellation is not rectangular. This paper proposes a simple geometrical procedure to derive an expression for the exact probability of symbol error for an *M*-ary AMPM system in the presence of AWGN.

While other researchers recently developed an expression for the exact probability of symbol error for an odd-bit M-ary AMPM system in the presence of AWGN [7], their method was based on Craig's approach [12], which requires evaluating complicated integrals and results were only shown for an 8-CAP/QAM system. On the other hand, our paper proposes utilizing the simple familiar geometrical approach, which is used in analyzing even-bit QAM constellations, to obtain the expression for the exact probability of symbol error as well as bit error of an odd-bit M-ary AMPM system in the presence of AWGN. Additionally, the derived expression in this paper is verified using MATLAB-based computer simulation for an odd-bit M-ary ($M=2^{2k+1}$) AMPM system. The results for different modulation orders are also contrasted against the upper-bound limits for the probability of symbol error of such a system.

This paper shows that geometrical approaches can be used to evaluate the probability of symbol error in signal constellations, where the shape of decision regions is irregular and more complicated than just a square or rectangle. This geometrical approach is based on the simple and familiar concept of calculating the error probability for two-symbol Pulse Amplitude Modulated (PAM) system using the Maximum Likelihood Ratio (MLR) [1] [8] as the basis for decision.

This paper is organized as follows. Section II provides a brief overview of odd-bit M-ary signal constellations for which the expression of symbol error probability is derived. The derivation of the exact expression for the probability of symbol error as well as bit error using the geometrical approach is detailed in section III. Section IV compares the derived expression with previous work, simulation results, and upper-bounded system limits. Finally, the paper is concluded in Section V.

II. 2^{2K+1} M-ARY AMPM SYSTEMS

Non-rectangular 2^{2k+1} M-ary AMPM signal constellations can be represented as two offset-overlapped 2^{2k} rectangular constellations. For example, Fig. 1 shows an 8-ary AMPM system, where its constellation is broken into two rectangular constellations. Generally, the construction of each rectangular constellation is based on two orthonormal basis functions φ_1, φ_2 given by

$$\varphi_1(t) = \sqrt{2/\varepsilon_g}\, g(t)\cos(2\pi f_c t) \tag{1}$$

$$\varphi_2(t) = -\sqrt{2/\varepsilon_g}\, g(t)\sin(2\pi f_c t) \tag{2}$$

where $g(t)$ is the symbol shaping pulse and ε_g is the energy of the pulse $g(t)$. f_c is the center frequency of the modulated signal. The symbols in each constellation can then be represented via the orthonormal functions as

$$s_{ni} = A_{mi}\varphi_1 + B_{mi}\varphi_2\,,\ n = 1, \ldots M/2, i = 1,2 \tag{3}$$

where s_{ni} is the n^{th} symbol in the i^{th} constellation. The symbols s_{ni} is composed of the combination of two levels A_{mi} and B_{mi}, in the $\varphi_1$and φ_2 directions, respectively. The above representation of the symbol s_{ni} can be expressed as a space vector given by

$$s_{ni} = [A_{mi} \quad B_{mi}\,] = \left[\sqrt{\varepsilon_g/2}\cdot a_{mi} \quad \sqrt{\varepsilon_g/2}\cdot b_{mi}\,\right] \tag{4}$$

where a_{mi} is the m^{th} level in the φ_1 direction for the i^{th} constellation and, similarly, b_{mi} is the m^{th} level in the φ_2 direction for the i^{th} constellation. The levels a_{mi} and b_{mi} are expressed as

$$a_{mi} = \left(4m - 1 - 2\sqrt{M/2}\right)d\ ; \quad 1 \le m \le \sqrt{M/2},\ i = 1, 2 \tag{5}$$

$$b_{mi} = \left(4m - 1 - 2\sqrt{M/2}\right)d\ ; \quad 1 \le m \le \sqrt{M/2},\ i = 1, 2 \tag{6}$$

where d is the distance between two consecutive a_{mi} or b_{mi} levels and $M = 2^{2k+1}$ is the total number of symbols in the M-ary AMPM system. Observe that the total number of symbols in the above 2-constellation system is $2\sqrt{M/2}^{\,2} = M$.

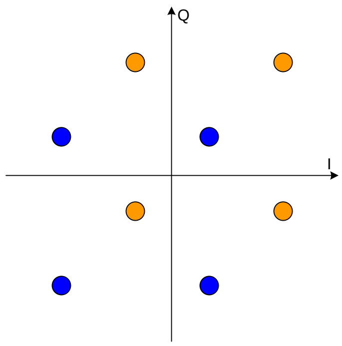

Fig. 1. Constellation of an 8-ary AMPM system is broken into two rectangular constellations

Calculating the average symbol energy of all symbols in the *M*-ary AMPM system, ε_s , consists of averaging the energy of all the equally likely symbols in both constellations as 5

$$\varepsilon_s = \frac{1}{M}\sum_{\substack{n=1\ldots\frac{M}{2}\\ i=1,2}} \varepsilon_{ni} = \frac{1}{M}\sum_{\substack{n=1\ldots\frac{M}{2}\\ i=1,2}} |s_{ni}|^2 \tag{7}$$

where ε_{ni} is the energy of the symbol s_{ni}. Expressing ε_s as the average energy of all symbols in *both constellations*, (7) can be rewritten as

$$\varepsilon_{av} = \frac{1}{M}\left[2 \cdot \sum_{m_1=1}^{\sqrt{\frac{M}{2}}} \sum_{m_2=1}^{\sqrt{\frac{M}{2}}} \left\{ \left(\sqrt{\frac{\varepsilon_g}{2}} \left(4m_1 - 1 - 2\sqrt{M/2}\right) d \right)^2 + \left(\sqrt{\frac{\varepsilon_g}{2}} \left(4m_2 - 1 - 2\sqrt{M/2}\right) d \right)^2 \right\} \right] \tag{8}$$

From (8), it can be shown that the average symbol energy in 2^{2k+1} *M*-ary AMPM system is given be

$$\varepsilon_s = \frac{\varepsilon_g}{3} d^2 (2M - 1) \tag{9}$$

The detector of the 2^{2k+1} *M*-ary AMPM system consists of two correlators that use φ_1 and φ_2 as reference signals. After integration over the symbol duration, the output of both correlators form coordinates of the received symbol. The process of symbols decoding depends on the MLR concept [8], where the ideal symbol closest to the received symbol (i.e., minimum Euclidean distance) is selected to be the output of the symbol detector. Since all symbols are equally likely, the decision regions represent the half-point traces between ideal symbols as shown in Fig. 2 and 3. Observe that the constellations in Fig. 2 and Fig. 3 contain irregular-shape decision regions, which are more complicated than the familiar rectangular decision regions found in rectangular QAM systems.

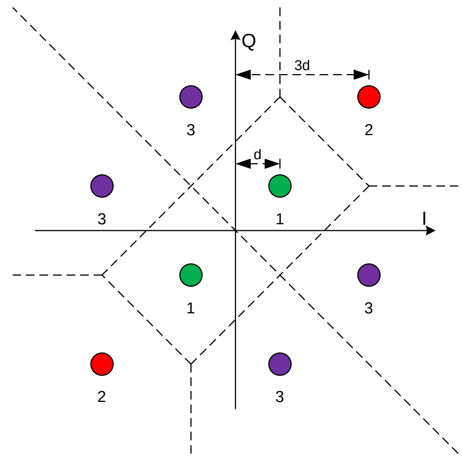

Fig. 2. Different decision region types composes the constellation of an 8-ary AMPM system

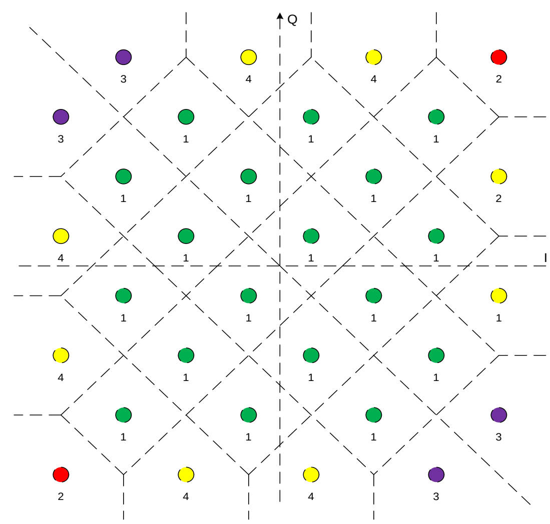

Fig. 3. Different decision region types composes the constellation of a 32-ary AMPM system (General case of 4 decision region types)

III. Derivation of the Error Performance Expressions

In this section, we utilize geometrical approach to develop closed-form expressions for the Symbol Error Probability (SEP) and Bit Error Probability (BEP) of 2^{2k+1} *M*-ary AMPM system in the presence of AWGN.

Observe from Fig. 2 and 3 that the constellation of an *M*-ary AMPM system contains four different types of decision regions denoted by Type 1, Type 2, Type 3, and Type 4 on Fig. 2 and 3. The number of constellation points that belong to those regions are denoted by N_1, N_2, N_3, and N_4, respectively. They are given in Table 1.

When applying the MLR concept (shortest Euclidean distance decision rule) to decide which symbol has been transmitted, correct symbol detection occurs if the noise is small enough to keep the received symbol within the decision region of the transmitted symbol. Since the constellations under study contain four different decision regions types, the probability of correctly decoding a particular symbol will depend on the shape of the decision region of the transmitted symbol.

TABLE 1. NUMBER OF CONSTELLATION POINTS IN DIFFERENT DECISION REGIONS TABLE TYPE STYLES

Decision Region	**Number of constellation points**
Type 1	$N_1 = 2\left(\sqrt{M/2} - 1\right)^2$
Type 2	$N_2 = 2$
Type 3	$N_3 = 4$
Type 4	$N_4 = \left(\sqrt{M/2} - 2\right)4$
Total	$N_1+N_2+N_3+N_4 = M$

Therefore, the process of calculating the probability of decoding symbols correctly, P_c, involves calculating the probability of correct symbol decoding in every decision region type, which can be expressed as

$$P_c = \sum_{i=1}^{M} P_c(C|s_i)P(s_i) \tag{10}$$

where s_i is the i^{th} symbol in the M-ary AMPM constellation, $P(s_i)$ is the probability of the symbol s_i, and $P_c(C|s_i)$ is the probability of correct decoding given s_i was transmitted. When all symbols are equally likely, the probability of correct symbol decoding can be represented as

$$P_c = \frac{1}{M}\sum_{i=1}^{M} P_c(C|s_i) \tag{11}$$

Observe that the $P_c(C|s_i)$ term is a function of the transmitted symbol and therefore depends on the decision region that corresponds to that symbol. Since all symbols that correspond to each decision region are equally likely and there are only four different decision regions, the probability of correct symbol decoding, P_c, in (11) can be rewritten as

$$P_c = \frac{1}{M}\left[N_1 \cdot P_{c1} + N_2 \cdot P_{c2} + N_3 \cdot P_{c3} + N_4 \cdot P_{c4}\right] \tag{12}$$

where P_{c1}, P_{c2}, P_{c3}, P_{c4} and are the probability of correct symbol decoding given that the transmitted symbol corresponds to decision region Type 1, Type 2, Type 3, and Type 4, respectively.

The derivation starts by evaluating the probability of correct symbol decoding given that the transmitted symbol belongs to decision region Type 1 shown in Fig. 4. Using the familiar MLR concept after rotating the decision region by 45°, P_{c1} can be expressed as [1] [8]

$$P_{c1} = P\left(|n_a| < \sqrt{2}d\sqrt{\frac{\varepsilon_g}{2}}, |n_b| < \sqrt{2}d\sqrt{\frac{\varepsilon_g}{2}}\right) \tag{13}$$

where n_a and n_b are the zero-mean noise components in the φ_1 and φ_2 directions, respectively, with variance of $N_o/2$, where $N_o/2$ represents the two-sided AWGN power spectral density level. It can be shown that (13) can be rewritten as [8]

$$\begin{aligned} P_{c1} &= \left(1 - 2P\left(|n_a| > \sqrt{2}d\sqrt{\frac{\varepsilon_g}{2}}\right)\right)\left(1 - 2P\left(|n_b| > \sqrt{2}d\sqrt{\frac{\varepsilon_g}{2}}\right)\right) \\ &= \left(1 - 2Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)\left(1 - 2Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right) \\ P_{c1} &= 1 - 4Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + 4Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \end{aligned} \tag{14}$$

where $Q(\cdot)$ is the familiar often-tabulated function, which is the tail probability of the standard normal distribution [1].

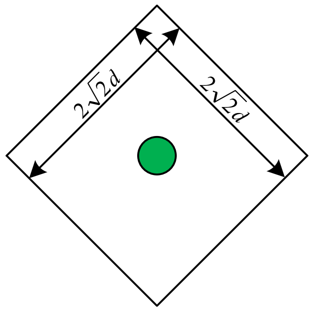

Fig. 4. Decision region Type 1 in an *M*-ary AMPM system

Next, P_{c2}, which represents the probability of correct symbol decoding given that the transmitted symbol corresponds to the more-complicated decision region Type 2, is evaluated using geometrical approaches similar to those developed in [11]. The decision region Type 2 represents the shaded area in in Fig. 5, where P_{c2} is equal to the area R_2' minus the area R_2'', where both areas are evaluated under the noise Gaussian distribution curve. That is, P_{c2} can be represented as

$$P_{c2} = P_c(C|s_2) = P_c(R_2|s_2) = P_c(R_2'|s_2) - P_c(R_2''|s_2)$$

$$P_c(R_2|s_2) = P_{R_2'} - P_{R_2''} \tag{15}$$

where s_2 implies that the transmitted symbol corresponds to decision region Type 2. The probability of correct decision occurring in R_2' (i.e., area of R_2'), is easily calculated using the MLR principle described above and is given by

$$P_{R_2'} = P_c(R_2'|s_2) = \left(1 - Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)\left(1 - Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)$$

$$P_{R_2'} = 1 - 2Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + Q^2\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{16}$$

In order to find P_{c2}, the area R_2'' still needs to be calculated. R_2'' can be geometrically found by taking one-quarter the difference of square areas *ABCD* and *WXYZ* shown in Fig. 6, which can be expressed as

$$P_{R_2''} = P_c(R_2''|s_2) = P_c(WBX|s_2) = P_{WBX} = \frac{1}{4}(P_{ABCD} - P_{WXYZ}) \tag{17}$$

where P_{WBX}, P_{ABCD}, and P_{WXYZ} are the areas of the *WBX* triangle, *ABCD* rectangle, and *WXYZ* rectangle, respectively, evaluated under the noise Gaussian distribution curve. P_{ABCD} can be easily found to be

$$P_{ABCD} = P_c(ABCD|s_2)$$

$$P_c(ABCD|s_2) = \left(1 - 2Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)\left(1 - 2Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)$$

$$P_c(ABCD|s_2) = \left(1 - 2Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)^2 \tag{18}$$

P_{WXYZ} can be expressed as

$$P_{WXYZ} = P_c(WXYZ|s_2)$$

$$P_c(WXYZ|s_2) = \left(1 - 2Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)^2 \tag{19}$$

Therefore, from (17), (18), and (19), can be expressed as

$$P_{R_2''} = P_{WBX} = -Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + Q^2\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) - Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{20}$$

Using (15), (16), and (20), the area P_{c2} can be expressed as

$$P_{c2} = 1 - Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) - Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{21}$$

Next, P_{c3}, which represents the probability of correct detection, given that the transmitted symbol corresponds to decision region Type 3, is evaluated. P_{c3}, which is equivalent to the shaded area (R_3) in Fig. 7, can be obtained by subtracting the triangle area R_3'' from the triangle area R_3' and then adding the rectangular area R_3''' to the difference, which can be written as

$$P_{c3} = P_c(C|s_3) = P_c(R_3|s_3)$$

$$P_c(R_3|s_3) = P_c(R_3'|s_3) - P_c(R_3''|s_3) + P_c(R_3'''|s_3)$$

$$P_{c3} = P_{R_3'} - P_{R_3''} + P_{R_3'''} \tag{22}$$

where s_3 implies that the transmitted symbol corresponds to decision region Type 3. Observe that $P_{R_3'}$ (the area of R_3') can be easily expressed as

$$P_{R_3'} = P_c(R_3'|s_3) = \frac{\left(1-Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)\left(1-Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)}{2}$$

$$P_{R_3'} = P_c(R_3'|s_3) = \frac{1-2Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)+Q^2\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)}{2} \tag{23}$$

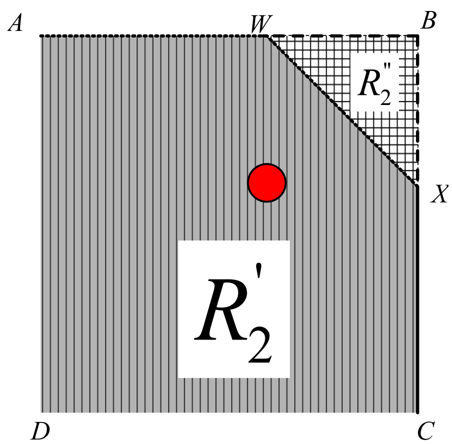

Fig. 5. Detailing decision region Type 2 (R_2): R_2 is the Gray-shaded area in the figure, which equals R_2' minus R_2'', where R_2' is the open-ended vertically-hashed square area (*ABCD*) and R_2'' is the horizontally-hashed triangular area (*WBX*).

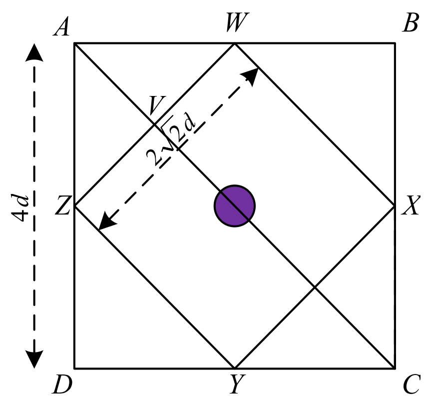

Fig. 6. Calculating the area R_2'' in Fig. 5. R_2'' is one-quarter the difference of square areas *ABCD* and *WXYZ*

It can be shown that the area R_3'', is half of the *WXB* triangle area calculated in (20). Therefore, the area R_3'' is given by

$$P_{R_3''} = P(R_3''|s_3) = P_c(AWV|s_2) = \frac{P_c(WBX|s_2)}{2}$$

$$P_{R_3''} = \frac{-Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)+Q^2\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)+Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)-Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)}{2} \tag{24}$$

Finally, the area R_3''' (*ZVCD*) is the half of the open-ended rectangular area (*ZWXD*) shown in Fig. 7, and therefore can be easily represented by the following equation

$$P_{R_3'''} = P_c(R_3'''|s_2) = \frac{\left(1-Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)\left(1-2Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)}{2}$$

$$P_{R_3'''} = P_c(R_3'''|s_2) = \frac{1-3Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)+2Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)}{2} \tag{25}$$

Using equations (22) through (25), it can be shown that P_{c3} is given by

$$P_{c3} = 1 - \frac{1}{2}Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) - 2Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + \frac{3}{2}Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{26}$$

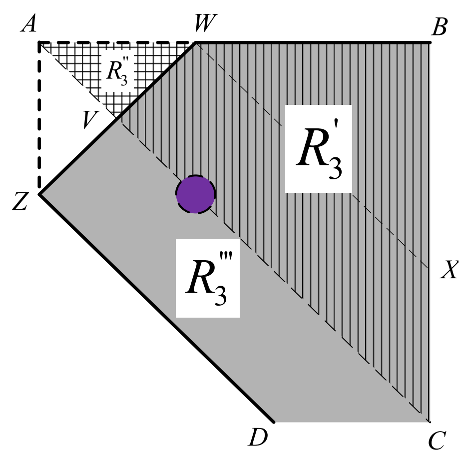

Fig.7. Detailing decision region Type 3 (R_3): R_3 is the Gray-shaded area in the figure, which equals $R_3' - R_3'' + R_3'''$, where R_3' is the open-ended vertically-hashed triangular area (*ABC*), R_3'' is the horizontally-hashed triangular area (*AWV*), and R_3''' is the open-ended rectangular area (*ZVCD*).

The last step in the analysis is to evaluate P_{c4}, which represent the probability of correct symbol detection given that the transmitted symbol corresponds to decision region Type 4. P_{c4}, which represents the shaded area R_4 in Fig. 8, can be obtained by subtracting the two triangles R_4'' and R_4''' from the rectangular area R_4', and therefore can be written as

$$P_{c4} = P_c(C|s_4) = P_c(R_4|s_4)$$

$$(C|s_4) = P_c(R_4|s_4) = P_c(R_4'|s_4) - P_c(R_4''|s_4) - P_c(R_4'''|s_4)$$

$$P_{c4} = P_{R_4'} - P_{R_4''} - P_{R_4'''} \tag{27}$$

where s_4 implies that the transmitted symbol corresponds to decision region Type 4. Using the MLR principle, it can be shown that the rectangular area R_4' is given by

$$P_{R_4'} = P_c(R_4'|s_4) =$$
$$\left(1 - Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)\left(1 - 2Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right)\right)$$
$$P_{R_4'} = 1 - 3Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + 2Q^2\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{28}$$

The triangular areas R_4'' and R_4''' are equal and can be found in a similar fashion to the area and therefore the area is given by (20), which is repeated here for convenience

$$P_{R_4''} = P(R_4''|s_4) = P_c(AWZ|s_4) = P_c(WBX|s_4) = P_{R_4'''}$$

$$P_{R_4''} = P_{R_4'''} = -Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + Q^2\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) - Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{29}$$

Using (27), (28), and (29), P_{c4} is evaluated and found to be

$$P_{c4} = 1 - Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) - 2Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + 2Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \tag{30}$$

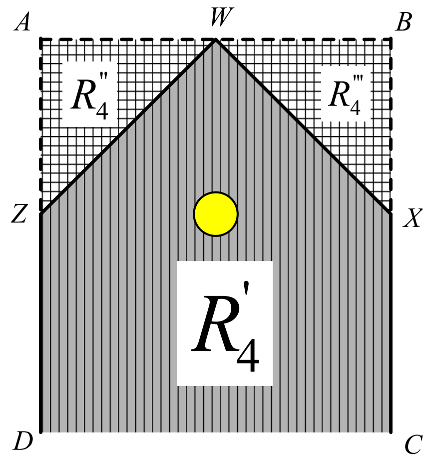

Fig. 8. Detailing decision region Type 4 (R_4): R_4 is the Gray-shaded area in the figure, which equals $R_4' - R_4'' - R_4'''$, where R_4' is the open-ended vertically-hashed rectangular area (*ABCD*), R_4'' is the horizontally-hashed triangular area (*AWZ*), and R_4''' is the horizontally-hashed triangular area (*WBX*).

Finally, the total probability of correct symbol decoding, P_c, is found using (12), (14), (21), (26), and the values in Table 1. After few algebraic simplification steps, the result is given in the following expression

$$\begin{aligned} P_c = 1 &+ \left(-4 + 8\sqrt{\frac{1}{2M}} - \frac{2}{M}\right)Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \\ &+ \left(4 - 8\sqrt{\frac{1}{2M}}\right)Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \\ &+ \left(-4\sqrt{\frac{1}{2M}} + \frac{4}{M}\right)Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \end{aligned} \tag{31}$$

The probability of symbol error for the 2^{2k+1} M-ary AMPM system in in presence of AWGN is $P_e = 1 - P_c$, which results in the following equation

$$\begin{aligned} P_e = &\left(4 - 8\sqrt{\frac{1}{2M}} + \frac{2}{M}\right)Q\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \\ &+ \left(8\sqrt{\frac{1}{2M}} - 4\right)Q^2\left(\sqrt{2}\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) + \left(4\sqrt{\frac{1}{2M}} - \frac{4}{M}\right)Q\left(2\sqrt{d^2\frac{\varepsilon_g}{N_0}}\right) \end{aligned} \tag{32}$$

Using (9), (32), and the fact that the symbol energy is r times the bit energy, where r is the number of bits per symbol ($r = log_2(M)$), P_e can be rewritten as

$$P_e = \left(4 - 8\sqrt{\frac{1}{2M}} + \frac{2}{M}\right) Q\left(\sqrt{\frac{6\log_2(M)}{(2M-1)}\frac{\varepsilon_b}{N_0}}\right)$$

$$+\left(8\sqrt{\frac{1}{2M}} - 4\right) Q^2\left(\sqrt{\frac{6\log_2(M)}{(2M-1)}\frac{\varepsilon_b}{N_0}}\right)$$

$$+\left(4\sqrt{\frac{1}{2M}} - \frac{4}{M}\right) Q\left(\sqrt{\frac{12\log_2(M)}{(2M-1)}\frac{\varepsilon_b}{N_0}}\right) \quad (33)$$

which coincides with the expression provided in [7], where Craig's approach, which requires evaluating complicated integrals, was used to obtain the probability of error expression. Observe that the above expression for P_e is valid for all Odd-bit M-ary AMPM systems. Therefore, it can be used to obtain the specific probability of symbol error expressions for different odd-bit AMPM QAM systems as follows

For $M = 8$, P_e is found to be

$$P_e = \frac{9}{4} Q\left(\sqrt{\frac{6}{5}\frac{\varepsilon_b}{N_0}}\right) - 2Q^2\left(\sqrt{\frac{6}{5}\frac{\varepsilon_b}{N_0}}\right) + \frac{1}{2} Q\left(\sqrt{\frac{12}{5}\frac{\varepsilon_b}{N_0}}\right) \quad (34)$$

While for $M = 32$, P_e is given by

$$P_e = \frac{49}{16} Q\left(\sqrt{\frac{10}{21}\frac{\varepsilon_b}{N_0}}\right) - 3Q^2\left(\sqrt{\frac{10}{21}\frac{\varepsilon_b}{N_0}}\right) + \frac{5}{8} Q\left(\sqrt{\frac{20}{21}\frac{\varepsilon_b}{N_0}}\right) \quad (35)$$

and for $M = 128$, P_e is expressed as

$$P_e = \frac{225}{64} Q\left(\sqrt{\frac{14}{85}\frac{\varepsilon_b}{N_0}}\right) - \frac{7}{2} Q^2\left(\sqrt{\frac{14}{85}\frac{\varepsilon_b}{N_0}}\right) + \frac{9}{32} Q\left(\sqrt{\frac{28}{85}\frac{\varepsilon_b}{N_0}}\right) \quad (36)$$

IV. Results – System Simulation and Comparison Against the Upper Bound Limit

MATLAB®-based system simulation was performed to validate the derived expressions. Figure 9 shows complete agreement between the theoretical expressions and simulation results for different modulation orders.

The upper bound approximation limit for an odd-bit M-ary ($M=2^{2k+1}$) AMPM system is given by [reference]

$$P_e < 4\left(1 - \frac{3}{2M}\right) Q\left(\sqrt{\frac{6\log_2(M)}{2M-1}\frac{\varepsilon_b}{N_0}}\right) \quad (37)$$

The theoretical expressions in (34) through (36) where compared against the corresponding upper-bound limits in Fig. 10. As expected, observe that the theoretical curves fall below the upper-bound limits.

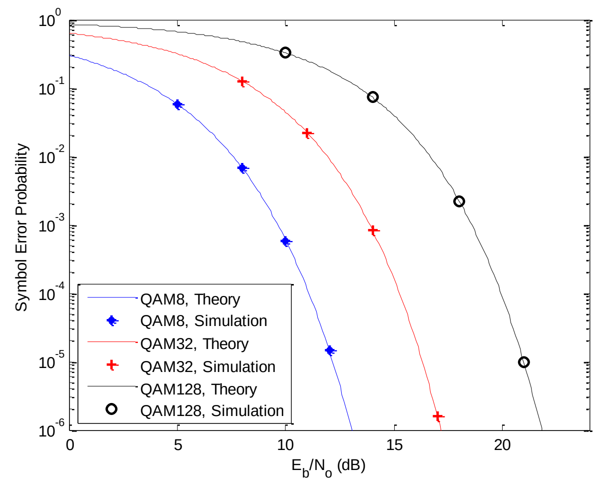

Fig. 9. Exact theoretical expression results match simulations results for different modulation orders of an odd-bit M-ary AMPM system.

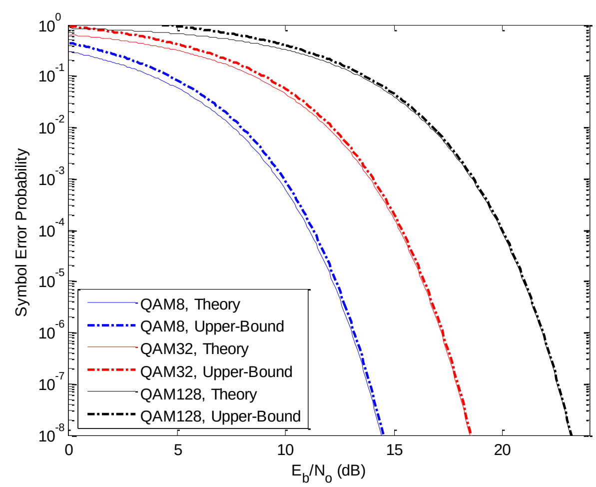

Fig. 10. Exact theoretical expression results, for different modulation orders of an odd-bit M-ary AMPM system, sit below the known system upper-bound limits.

V. CONCLUSIONS

This paper proposed the use of a simple geometrical approach to develop an expression for the probability of symbol error of an odd-bit M-ary ($M=2^{2k+1}$) AMPM system in the presence of AWGN. The obtained theoretical expressions were validated using computer-based system simulations and were compared against the known upper bound limits for such systems.

REFERENCES

[1] G. Proakis, *Digital Communications*, 4th Ed., McGraw-Hill Inc., 2001.

[2] Cable Television Laboratories, Inc., *Radio Frequency Interface Specification (CM-SP-RFIv2.0-C02-090422), Data-Over-Cable Service Interface Specifications (DOCSIS 2.0)*, 2009.

[3] Cable Television Laboratories, Inc., *Physical Layer Specification (CM-SP-PHYv3.0-I09-101008), Data-Over-Cable Service Interface Specifications (DOCSIS 3.0)*, 2010.

[4] J. Cioffi, T. Starr, M. Sorbara, P. J. Silverman, S. Thomas, *DSL Advances*, Prentice Hall PTR, 2002.

[5] M.G. Shayesteh, "Exact symbol and bit error probabilities of linearly modulated signals with maximum ratio combining diversity in frequency nonselective Rician and Rayleigh fading channels," *IET Commun*, Vol. 5, Iss. 1, pp. 12–26, 2011.

[6] P. Fines, A.H. Aghvami, "Performance evaluation of high level coded modulation over satellite channels," *in IEEE GLOBECOM*, pp. 417-421, 1992.

[7] M. Vaezi, J. Habibi Markani, "Exact Expression and a Simple Tight Upper Bound for the SER of Odd CAP/QAM Constellation," *in IEEE 68th Vehicular Technology Conference*, pp. 1-4, 2008.

[8] B. Sklar, *Digital communications*, 2nd Ed., Prentice Hall, 2001.

[9] M. Naeem, S. S. Shah, H. Jamal, "Performance Analysis of Odd Bit QAM Constellation," *in IEEE Symposium on Emerging Technologies*, pp. 178 – 181, 2005.

[10] H. Xu, "Symbol Error Probability for Generalized Selection Combining Reception of M-QAM," *in SAIEE Africa Research Journal*, Vol. 100, No. 3, pp. 68-71, 2009.

[11] K.H. Li, P.Y. Kam, "On the Error Performance of 8-QAM Involving a π/2-Wedge-Shaped Decision Region," *Communications and Signal Processing, 2003 and the Fourth Pacific Rim Conference on Multimedia. Proceedings of the 2003 Joint Conference of the Fourth International Conference on Information*, vol. 2, pp. 884-887, 2003.

[12] J. W. Craig, "A new, simple and exact result for calculating the probability of error for two-dimensional signal constellations," in *IEEE MILCOM'91Conf. Rec.*, Boston, MA, 1991, pp. 571-575.

A Software Defined Networking Approach To Cable Wi-Fi

Alon Bernstein, Vince Pandolfi, Charles Duffy, Sangeeta Ramakrishnan
Cisco Systems

Abstract

Software Defined Networking (SDN) is a framework that lends itself to coordinating different networks. Because of that it has the potential to tie Wi-Fi, the HFC and the core/access in a consistent and manageable way. The paper will present several options for this coordination – from evolutionary to game changing.

SDN definition?

Software Defined Networking (SDN) is a term that started being used in the networking industry around 2009. SDN can trace it's lineage to the Clean Slate Research Program at Stanford University [1] whose mission was to explore what kind of Internet we would design if we were to start with a clean slate and 20-30 years of hindsight. SDN is defined as the separation of Control and Data planes using an open standard protocol to communicate between them as depicted in Figure 1. This differs from a traditional network device (such as a Router, Switch, or CMTS) in which the data and control planes are vertically integrated. SDN promises flexibility and rapid innovation by virtue of the fact that the control software would be removed from the relatively constrained network device to a generic server that can be easily scaled to have more processing and memory capabilities.

In the SDN architecture, key functions such as routing, topology discovery, and policy are removed from the individual devices and located centrally in a controller and the applications that reside on top of it. The controller, with its " bird's eye" view of the network uses a simple protocol such as Openflow [2] to populate the forwarding tables of the networking devices. Figure 2 illustrates the Classic network architecture vs. SDN.

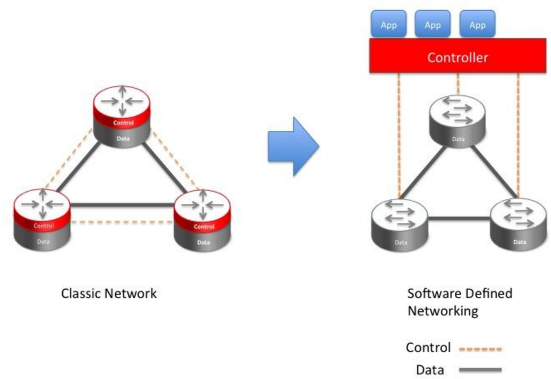

Figure 1 centralized vs. distributed control

In the center of the diagram is a Controller that handles communications to network devices including physical and virtual switches via a southbound protocol like OpenFlow. Northbound, the controller presents a level of abstraction of the underlying network. Applications communicate with the controller using "controller APIs" and the controller in turn interacts with the network. The applications could be written by the vendor of the controller, by any third party or by a service provider. In other words the controller acts as middleware that helps in providing a higher lever of abstraction to the application developers.
Anouther key attribute of SDN is the two-way communication with the network devices. The network can be thought of as a large distributed database of flows and states. Current configuration tools tend to either uni-directional (CLI which does not have a good error reporting) or localized to only one device (PCMM which is limited to the CMTS) limiting the ability of the configuration tools to rollback when an error occurs or verify the end-to-end correctness of a configuration.

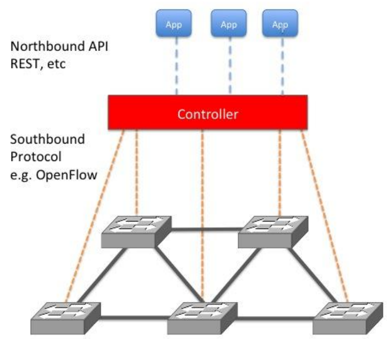

Figure 2 SDN architecture

It is worth calling attention to three emerging SDN related technologies. These are Programmatic APIs, Overlay Networks, and Network Function Virtualization (NFV). The first technology is where open published APIs are provided for existing and new network devices. These APIs allow an operator to control many if not all of the functions of the device, beyond the basic packet forwarding control that is provided by openflow. For example, an APIs could allow an application to change routing or implement QOS policy or simply retrieve the network topology. The second technology is Overlay Networks. The concept here is similar to a Virtual Machine where software creates an "overlay" network that rides on top of an "underlay" or traditional network. The Overlay replicates the functionality provided by a traditional network but since it is completely defined in software it is potentially more flexible. The overlay relies on the underlay network to provide stable transport and connectivity using all of the methods and protocols found in traditional networks. The third technology, NFV is being developed by the European Telecommunications Standards Institute (ETSI) in the NFV Industry Specification Group. NFVs goal is virtualize network functions such as Network Address Translation (NAT), Session Border Controllers (SBC) Deep Packet Inspection (DPI), etc. by using industry standard server virtualization techniques. In contrast to SDN, which moves the control plane to the cloud, NfV moves the data plane as well.

Matching SDN framework capabilities to Wi-Fi requirements

SDN is a framework that can apply to a large range of application. Part of the challenge of applying SDN is a good "divide and conquer" strategy. To decide what to focus on we will list the requirements for supporting a Wi-Fi network and then match them up with the SDN framework capabilities.

To support Wi-Fi networks the following requirements need to be met:

Mobility: Including micro-mobility, hostspots mobility and mobile offload (which will be explained in more detail later in the document).
Authentication: Validate subscriber identity
Subscriber Services: QoS, BW accounting, legal intercept, charging, parental control etc.
Access point (AP) management: SW upgrade, configuration
Mobile Node (MN) management: capability discovery etc.
RF resource management: Frequency planning
Seamless handoff: minimize packet drops when moving between access points.

Though SDN is a framework that may be defined in several ways, the following list is a high level set of capabilities that should fit all SDN flavors along with the benefit of each capability:

Separation of Control from Data plane: Feature velocity, easier debug, and fault tolerance
Logically centralized control plane: better control over the network, feature velocity, easier Debug/test/simulation.
Control plane can run off-box (aka "cloud"): developing SW in a non-embedded

environment, hope for better scale, and access to modern SW development tools
Use openflow or higher layer formal API: solve vendor interoperability issues, "end of protocols" (since a protocol does not have to be defined a-priority for two endpoints to interoperate), Fast release of new features; ISP can do its own SW customization
Create dedicated paths through the network: Assure services, simplify the network by flattening
Direct application control of the network: simplified configuration (auto-configuration/automation).

This paper focuses on the intersection of <u>mobility</u> requirements and <u>SDN capabilities</u> and outlines where SDN can help with improving mobility solutions.

<u>PMIPv6, SoftGRE and anchor points</u>

Since the concept of an anchor point is essential to our discussion we will start by examining how it applies to two common mobile arch1itectures before diving into the SDN discussion:

1. PMIPv6
2. SoftGRE

PMIPv6 is an IETF standard (RFC5213[4] and RFC5844), and provides mobility to endpoints, without requiring client modifications. PMIPv6 involves Mobility Access Gateway (MAG) and Local Mobility Anchor (LMA). LMA is defined to be the topological anchor point i.e. home agent for the Mobile Node's (e.g. Wi-Fi user device's) IP prefix (es) and manages MN's binding state via MAG. The MAG manages mobility-related signaling for the MN that is attached to its access link and is responsible for tracking the MN's movements to and from the access link and for signaling to the LMA.

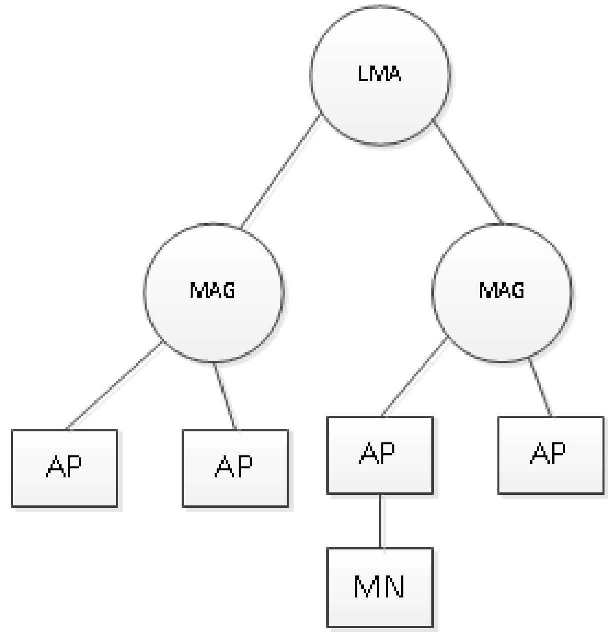

Figure 3 PMIPv6 Components

In PMIPv6, the MAG informs the LMA about the Wi-Fi users during user authentication/authorization. This allows the MAG and LMA to send/receive Wi-Fi user traffic (i.e. Ethernet frames) over the PMIPv6 tunnel. The MAG functionality could be embedded in a Wi-Fi Access Point, or on the CMTS. While the Wi-Fi user is connected to AP/MAG at layer2, its IP address is anchored the LMA. This allows IP mobility, when the Wi-Fi user roams and changes AP/MAG attachments.

<u>SoftGRE</u>

The architectural approach with softGRE is to build an over-the-top IP tunnel to deliver the Wi-Fi user device's Ethernet traffic between AP and a remotely located anchor point (i.e. tunnel termination entity), using GRE. This approach requires IP connectivity between AP and the centralized entity. The data plane comprises users' "Ethernet over GRE over IPv4|v6 over Ethernet [over DOCSIS (or PON)]" in the last-mile access and "Ethernet over GRE over IPv4|v6" (over MPLS, if existed) in rest of the network (up to that centralized entity). While Ethernet over GRE over IP usage is not well known or used, it is standardized at the IETF [RFC1771].

With softGRE an AP establishes a L2TP tunnel with the remote L2TP tunnel concentrator (e.g. centralized entity) and sends/receives Wi-Fi user device's Ethernet frames, over GRE (over IP) tunnel. It is important to note that GRE doesn't require a control channel and can be set up in a stateless manner (aka soft GRE) without requiring any tunnel configuration.

Although both approaches described above use an anchor point, there are differences between the two. For further details on the trade-offs of these architectures refer to [3].

For the reminder of the paper we will refer to an "anchor" without detailing the protocol framework around it.

How can SDN help?

There are several benefits for looking at SDN to solve challenges in SP Wi-Fi deployments. The first is that Wi-Fi mobility is in its infancy. The L1/L2 micro roaming has been efficient and enhanced by several standards like 802.11i/r/u. The greenfield exists in the area of macro roaming. The CAPWAP RFC5415 (see ref [7])_standard provides some limited geographic mobility and the concept of a separated control and data plane but the current implementation by vendors generally co-locates processing of both the control and data plane. The CAPWAP protocol provides a layer 2 transport and requires handoff of data traffic to other services like Internet services gateways and network address translation.

This leaves the service providers with having to stitch the network paths together for these different services. The use of the overlay tunnels allows for over the top deployments that may or may not be traffic engineered correctly into the core network. This is a particular challenge with larger providers that have silos of responsibility in their organizations.

There are numerous moving parts in a successful SP Wi-Fi implementation. The challenge becomes connecting these moving parts within and in many cases between services provides. The number of these touch points increases when considering the numerous use cases for SP Wi-Fi.

The use of SDN provides a dynamic network fabric to connect these different elements. The benefit increases exponentially when these components are provided by multi-vendor solutions.

There are two types of solutions SDN can offer:

1. **Evolutionary:** current Wi-Fi solutions rely on an anchor point in one way or another, i.e. Even though the MN moves between AP's the anchor remains a static point of connection to the rest of the network. An example of an anchor point of PMIPv6 would be the LMA. For softGRE it would be the wireless LAN controller. For anchored architectures we can identify current issues and outline how SDN may help address them.
2. **Game changing:** with SDN we create a Wi-Fi architecture without an anchor point. Instead of anchoring a MN to a fixed point so that the rest of the network can treat it as a fixed asset SDN allows the network (or more precisely a network overlay) to move along with the subscriber.

Application of SDN to anchored architectures

Mobility solutions range in complexity depending on how far from the home network a subscriber can get, but all the Wi-Fi anchored architecture share the same base concept: they maintain IP address persistency by tunneling traffic from the MN, so even though the tunnel end points may move as the

MN moves, the MN still “feels” as if its connected to its home network. Different MSO network implement (or considering implementing) architectures from basic IP persistence to mobile offload. This section will go into the details of each one of these solutions and how SDN may help with the mobility aspects:

1. **Basic IP address persistence within a domain**: The simplest mobility solution is to have a single anchor point – a direct tunneling of SSID traffic to a fixed anchor point. This means that subscribers can maintain IP persistency only within a “domain’ (which in reality can be fairly large). One simple example of this approach is an Intra-CMTS mobility mechanism whereby an MN keeps the same IP address as long as it moves between Access Points that are connected to the same CMTS. The current issues with the this solution are (a) when a handoff is needed a new tunnel setup needs to be established with a new anchor point which slows down the handoff process, (b) a single anchor point creates a large failure domain. On the positive side these tunneling/anchoring solution to IP persistence are widely deployed, relatively vendor neutral and simple. The way SDN may help in this architecture is to add the flexibility to attach to multiple anchor points by dynamically moving the tunnels without the need for the AP to establish a new one.
2. **Intra MSO mobility**: a step up from case (1) is IP address persistence anywhere within the MSO network (i.e. across domains and anchor points). Current solutions such as softGRE and PMIPv6 allow for this mobility and solve many of the problems with case (1) by having the ability to deal with MN movement across anchor points. The main help SDN can offer in this architecture is dynamic configuration of the anchor points and helping to decide on the optimal anchor point to connect to.
3. **Inter MSO mobility (branded services):** A subscriber may roam between MSO A and MSO B networks. From a technical implementation point of view the solution to this type of roaming is the same as case (2) described above and the main challenge is coordinating authentication and identity information between two different providers, which is outside the scope of this paper. From an SDN point of view this may require an even wider “bird’s view” of the network since the connection point to an external network have to be taken into account.
4. **Mobile offload**: mobile operators are interested in offloading their cellular data network to Wi-Fi hotspots, including MSO Wi-Fi hotspot. This would be the ultimate mobility play where the only common equipment is the MN itself. However, from a mobility point of view its still the same basic tunnel architecture as in the inter-MSO case (3) which in turn has the same data plane architecture as (2) and just like the previous case the challenges of exchanging authentication and identity across providers is out of scope for this paper. Note that with mobile offload the handoffs between the networks may be more frequent then in other use-cases because of the greater chance of overlap between the networks (mobile and WiFi).
5. **Local Breakout (application based)**: A mobile service can be finer-grained then MN mobility. Each application running on an MN might have its own mobility domain, therefor a single MN may require multiple anchor points to

connect to. For example, a user may have a dedicated connection to a home service provider for watching premium video content and at the same time, on the same MN, another connection for non premium content that runs in the background (messaging, e-mail etc.) and can be sent through the hosting MSO as a native service.

The use of tunnel technologies provides the mechanism to allow mobility in a geographic Wi-Fi deployment but the tunnels generally require termination on the anchor or tunnel endpoint. This behavior allows for seamless mobility but requires more extensive anchors due to higher amounts of signaling traffic and needing to service all user traffic – a local breakout based on applications puts further scaling requirements on an anchored system.

There are several popular methods to break traffic out of the tunnels. The first is Selected IP Traffic Offload (SIPTO see ref[6]). The SIPTO method supports offload of IP traffic directly to the Internet or other locally hosted services. This provides the best user network performance and prevents excessive traffic from needing tunnel transport back the anchor. One of the drawbacks would be loss of visibility of the user traffic flow for applications such as lawful intercept for the home provider and mobility for those applications that are offloaded. One of the other methods would be a concept of LMA (local mobility anchor) chaining. This allows Mobile IP traffic to be relayed between LMA anchor points based on the home relationship of the end user. This method does not necessarily provide for local breakout but larger provides could use this method for regional breakout of user traffic. This method does allow for home provider visibility and control, mobility and distrubted signaling load. SDN can help with dynamically selecting an LMA, which would be a simpler solution then LMA chaining. Note that with SDN intermediate nodes can do their own classification of the traffic and do not rely on the MN to tag application traffic.

For all the anchored architectures it becomes clear that although the control part might be unique the data plane part similar and that's not surprising given that the are all based on the premise of a tunnel to an anchor point. The SDN contribution can be summarized for all these technologies as:

1. **Optimized path selection**: when there are multiple anchor points to choose from an SDN solution has the advantage of seeing the network as a whole and dynamically choosing the best anchor based on various criteria (geographical proximity, load, cost and more)
2. **Dynamic anchor point selection**: in addition to selecting the optimal anchor point an SDN solution can assist in dynamically moving a tunnel to a new anchor point as network conditions or MN location change.
3. **SDN may address vendor interoperability issues**: different providers use different equipment and different tunneling technologies. However, for SDN enabled equipment we can configure a consistent tunnel across equipment and provider sites. Note that for the cases of inter-MSO mobility and mobile offload this will require controller-to-controller coordination, which is outside the scope of this paper.

It is worth mentioning that there are other mobile technologies, most notably LISP-MN that requires changing to the MN itself. LISP-

MN requires tunnel termination that is not strictly an "anchor". However, the discussion of LISP is outside the scope of this paper.

Un-anchored Architectures

With an un-anchored architecture there is no need to tunnel user traffic. The network can follow the subscriber and present a virtual port that appears constant to the MN even though it's physically moving. For example, the response to ARP/DHCP and other L2/L3 network protocol can mimic that of the original home port. It can be view as a connection to an MN that moves as the MN moves.

In the data center world a solution such as mentioned above already exist and is generally referred to as "network virtualization", though one should note that for the Wi-Fi case we discuss the primary use-case is not to create a whole network overlay, but rather to create a dedicated overlay to a particular user.

In its pure form the un-anchored architecture can become a scaling challenge for the network because every subscriber, and in fact, every application within an MN, will need its own dedicated "flow" or connection through the network. While network devices of the future might support such scale we propose an interim solution; all flows that share a common destination, for example a path outside the MSO and into the global internet (at some exit point which is coordinated by SDN as well) can share a tag and do not need to be managed as individual flows, thereby reducing the need to have per-flow scaling in the core of the network. Note that the aggregation point does not create a new anchor – it has nothing to do with mobility. Its role is only to aggregate flows in order to improve network scaling.

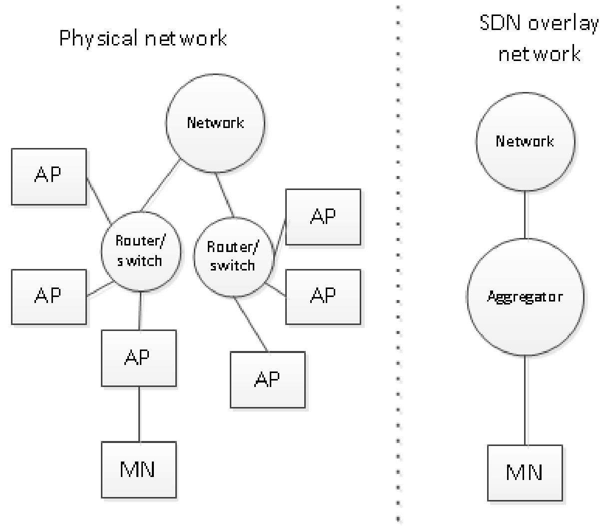

Figure 4 Overlay network with SDN

Figure 4 illustrates the differences between the overlay network and the physical network. The MN is logically connected to the virtual network (since it's the virtual network that responses to ARP/DHCP etc from the MN) and that network is a simple static network that does not change even as the MN moves between AP.

While one may argue that a tunnel is also a way to implement an overlay network, however, the SDN approach has several advantages; there is no need for a tunnel and therefor no MTU issues with a tunnel encapsulation or scaling issues with tunnel control. In addition the centralized control of the overlay allows for better optimization of network paths, and faster response to changing network conditions then a traditional tunnel.

Conclusion

SP Wi-Fi deployments are accelerating throughout the industry. Furthermore, there are elements in existing Wi-Fi protocols that can be considered "SDN" such as the separation of data and control in CAPWAP. Current deployments are able to provide a number of capabilities including nomadic access, mobility, and roaming capability.. However the typical SP Wi-Fi solution is fairly complex. Most implementations are based on some type of tunnel architecture from the access point to a controller/BNG. There may also be tunnels from these controllers to roaming partners and or managed services customers. These various over the top tunnels are generally implemented by several groups: a Wi-Fi group, network engineering, server operations, security and partner providers, all operating as independent silos. While existing architectures are functional and deployable, the use of SDN may simplify the architecture significantly. The use of SDN may allow a standard interoperable pipe between these service points, and more specifically, as covered in this paper, a more optimized and dynamic selection of anchor points with current architectures as well as complete elimination of anchor points.

References

1. http://cleanslate.stanford.edu
2. https://www.opennetworking.org/sdn-resources/onf-specifications/openflow
3. "Architectural approaches for integrating SP Wi-Fi in Cable MSO networks", Rajiv Asati, Sangeeta Ramakrishnan, Rajesh Pazhyannur, SCTE 2012.
4. "Proxy Mobile IPv6", http://www.ietf.org/rfc/rfc5213.txt
5. "Challenges of Unlicensed Wi-Fi Deployments: A Practical Guide for Cable Operators" SCTE 2012
6. Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO) 3GPP TR 23.829
7. Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification IETF RFC 5415

ACCOUNTING FOR TECHIES: TAKING IT TO THE ULTRA

Dr. Robert L. Howald
Dr. Sean T. McCarthy
ARRIS

Abstract

High-Definition video has been a terrific success story for cable operators. It is irresistible to consumers once they're introduced to the vividness of the HD experience. As a high-demand premium service, it delivers precious bottom-line value to the MSO. However, as more subscribers flock to HD, pressure increases on operators to provide more HD programming. A key dilemma becomes allocating sufficient bandwidth and balancing HD demands with competing demand for spectrum from IP services. Fortunately, all operators are engaged in multiple activities that will substantially improve network efficiency. Over time, these various tools allow operators to see a way out of this resource limitation predicament……or do they?

High-Definition service today is HD 1.0. Already in labs, standards bodies, and in demonstrations from trade show floors to "big box" stores are emerging HD 2.0 and even HD 3.0 technologies. The operator's spectrum efficiency toolkit is not improving in isolation. Technology to further enhance the media experience is occurring in parallel, and it comes at the price of more bits-per-second. Not far around the next corner, Quad Format HD (aka 4k HD) aims to provide the next phase of display resolution. High Frame Rate (HFR) HD is poised to make subscribers think differently about dynamic resolution. Finally, Ultra High-Definition, an HD 3.0 candidate, represents a 32x resolution experience compared to HD 1.0.

In this paper, we raid the engineering labs and turn over their technical reports to the accountants to develop the long term Balance Sheet. We will account for the resource Liabilities of increased pixels per frame and increased frames per second. On the Asset side of the ledger, we analyze how new formats interact with available knobs and levers in MPEG-4 (H.264), anticipate the emergence of HEVC (H.265), consider architecture and spectrum evolution, and discuss the numerical implications of each. We also integrate IP Video and assess its role in the network transport transition, in addition to the service transition. The service transition is encumbered by quantifiable Liabilities brought about by the nature and burden of legacy support and multi-format simulcast. We assimilate compounded IP growth rates as part of a full IP transition. Using our itemized analysis, we reconcile the Balance Sheet, and draw conclusions in the context of HFC capacity constraints and optimization. Finally, we present strategies for the HD 2.0 era and consider an era of HD 3.0.

INTRODUCTION

Decades of service evolution – video, voice, and data alike – have given operators a sound historical basis for business planning of new growth scenarios. The prevailing MSO approach has been a very successful pay-as-you grow approach, capitalizing on technologies as they mature while meeting consumer service demands. This approach has benefited from the availability of new HFC capacity, incrementally exploited through fiber deep extensions, RF bandwidth upgrades in the distribution plant, use of WDM to fuel continued segmentation and service options, and migrating to either all-digital services and/or switched digital video (SDV) architectures.

However, while the appetite for all things HD continues to be strong, the lifecycle of already-bandwidth-burdensome HD itself has only just begun. Cable systems deliver 720p and 1080i formats today, while "Full HD" – a term that never had a uniformly understood meaning itself – 1080p already exists in the consumer electronics world of Blu-Ray and gaming consoles. Flat panel televisions continue to become larger, more capable, and lower cost. The CES show earlier this year had dozens of "4K HD" (3840 x 2160p), televisions on display, and sizes as large as 84" (7 feet!) are on the market. 4K HD television sales have surprised analysts, and they have upwardly adjusted their forecasts. By 2015, it is now projected that the majority of TVs sold, if not the vast majority, will be 4K HD capable. Since these TVs are likely to come with format upconversion, users can take advantage of them before there is mass content delivered in the 4K format.

Another accelerating factor is the cycle of television replacement. It has dropped to about 6 years more recently, down from a historical 10 year lifecycle. The move from HD to Ultra High Definition (UHD – to be defined) is anticipated to not be as large of a hurdle for consumers as the move from SD to HD was. This is somewhat ironic, since the relative enhancement to viewing performance from SD to HD is larger than HD to 4K.

The 4K HD format represents 4x the number of pixels of 1080-column HD. Formats definition and early technology also exist beyond 4K UHD in the form of 8K UHD (aka NHK's Japan's Super Hi-Vision). The pixel multiplier for 8K HD is another 4x over 4K HD, with the total pixel grid standing at 7680 x 4320p. As always, "p" stands for progressive scan, but the implicit "60" cannot necessarily be assumed as new video science takes place, as we shall discuss.

With IP CAGR's racing ahead, and quantifiable limits of the HFC architecture's capacity, it is important to account for these emerging video formats destined for the marketplace when we analyze long-term capacity management. In the case of 4K UHD at least, with televisions moving off the shelves more quickly than expected, it suggests a high likelihood of a mass-market technology in a way that the cumbersome size implications of 8K UHD may not.

In the analysis to follow, we will do the accounting – Bandwidth Assets and Bandwidth Liabilities – that give us insight into the future possibilities. Then, because of the limitations of this simplistic Balance Sheet approach, we will build out a long term service and architecture Capacity Management Timeline. A key component of the timeline is that the service evolution will occur in parallel with a major architecture evolution – the IP Transformation. As we shall see, it is valuable – critical – to understand this end objective and the implications. The path to Network Nirvana is a complex balance of new services arriving, old services phasing out, and architectural techniques and tools leveraged adroitly to survive the capacity management challenge of this multi-dimensional transition.

HD 2.0 & HD 3.0: WE'VE ONLY JUST BEGUN

Pixel Perfect

Last fall (October 2012) the Consumer Electronics Association (CEA) came out with definitions of Ultra High Definition to help to market higher resolution televisions. And, it is not just the big box stores and CE vendors that are in on the act – the ITU is developing recommendations for both 4K HD and 8K HD. And, the European

Broadcasting Union now defines 4K HD and 8K HD as UHD-1 and UHD-2, respectively. We will use UHD-1 to refer to 4K and UHD-2 to refer to 8K when it is not explicitly stated throughout the paper. Depending on discussion context, we will also refer to UHD-1 as HD 2.0 and UHD-2 as HD 3.0.

Since the advent of HD, the video and CE industries have gained a strong understanding of the relationship among resolution, screen size, and viewing distance. Video technology has continued to advance as the science of video perception and factors effecting Quality of Experience (QoE) are better understood. The science now spans the knowledge base created through the development of HD, as well as human vision biology and the neural processing that interprets the visual information sent to the brain. Basically, the science now incorporates the entirety of the Human Visual System (HVS).

Beginning with the fundamentals of the interplay among screen size, resolution, and viewing distance, Figure 1 captures the interaction in a straightforward way [23]. For a fixed screen size, higher resolutions are best viewed by sitting closer to allow for the full benefit of the increased detail on the display. For a fixed distance from the display, the benefit of higher format resolutions ensues with a larger screen size. These are well-understood principles.

As a simple example from Figure 1, if viewing a 50" screen from more than 20 ft away or greater, you will lose the benefit of HD at 720p and instead have an experience more akin to Standard Definition 480p. Sitting too close, such as 5 ft away on a 100" 1080p screen, threatens quality due to the distinguishing of pixels. This actually explains UHD-2 (8K HD) resolution. This format was originally envisioned as an immersive experience using a wider field of view (FOV), and to achieve that FOV by sitting closer demanded very high resolution to ensure high image quality.

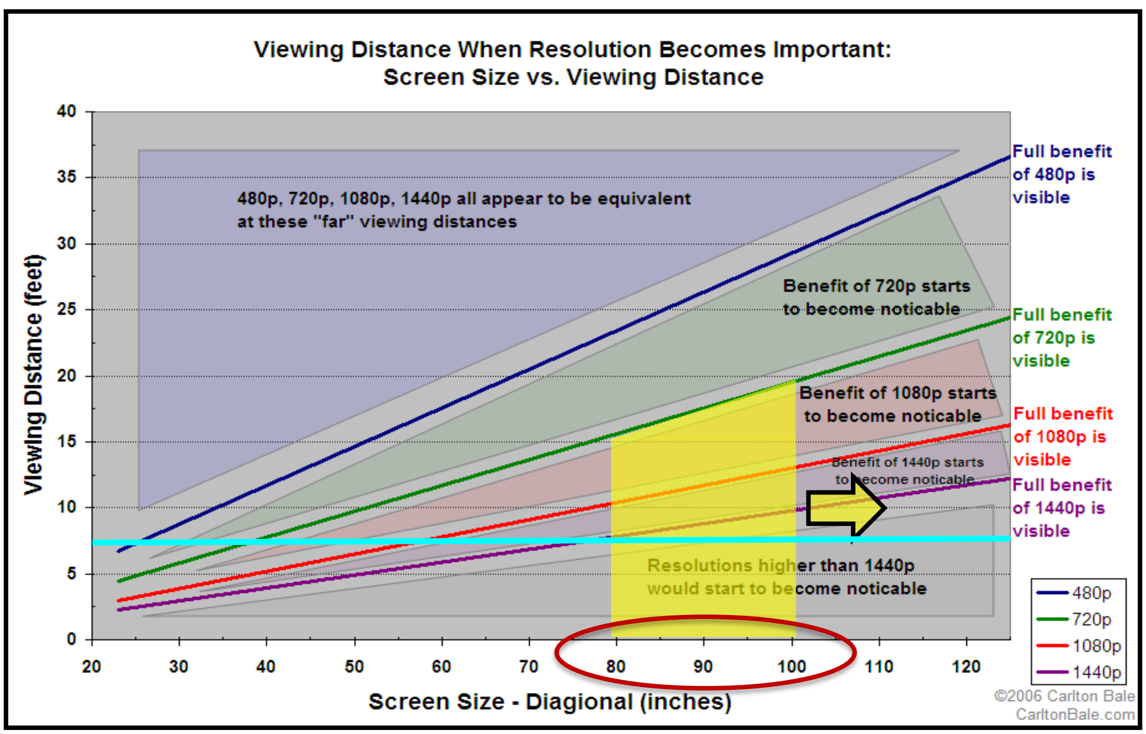

Figure 1 – Screen Size, Viewing Distance, and Spatial Resolution [23]

Consider the bottom right corner of Figure 1, shaded yellow. A typical viewing distance in the home today is about 7.5-9 feet. Flat panel screens are available now at ever-increasing sizes, such as those shown in this shaded yellow range. At 7.5 feet distance (light blue line), "only" a 55" screen could show perceptible benefits for resolutions better than 1080p (light blue line crosses red line). A 60" screen is sometimes considered the 4K TV threshold of benefit. Now 4K HD capable 84" flat panels are available (for a cool $25,000).

The Eyes Don't Have It

UHD recommendations are also correlated to visual acuity principles as they relate to the ability to resolve the image detail beyond simple optical acuity principles used to characterize the quality of our eyesight. This is an area where significantly better understanding has occurred in recent years. The common Snellen optics standard of 20/20 vision is a figure of merit based on visual acuity (VA) associated with accurately resolving high contrast, sharp edged objects (letters). However, it is well understood, for example, that the human visual system can perceive vernier or edge alignment at up to *ten times* the precision of standard VA, and such aspects of the brain's processing of video images translate to QoE.

Snellen acuity for "normal" vision works out to 30 cycles per degree. By comparison, as shown in Figure 2, HDTV satisfies a resolution consistent with 60 cycles per degree resolution, achieving "simple acuity," while UHD-1 achieves 120 cycles per degree, or the noticeable subjective threshold of "hyperacuity." This threshold accounts for the entire visual system's role in the viewing experience, beyond simple spatial resolution and vision receptor biology. It accounts for the role of neural processing. This experience achieves what's often referred to as "retinal" image quality [11].

Figure 2 – Tiers of Acuity Relate to the Quality of the Viewing Experience [11]

Simply put, treating the biology of the eye as a camera of a particular resolution and comparing it to image pixel density does not capture the essence of the brain's role in processing the information contained in the image, and which subjectively effecting our perception of its "realness." Indeed, studies to determine this relationship have concluded precisely this. An example of such a study is shown in Figure 3 [11].

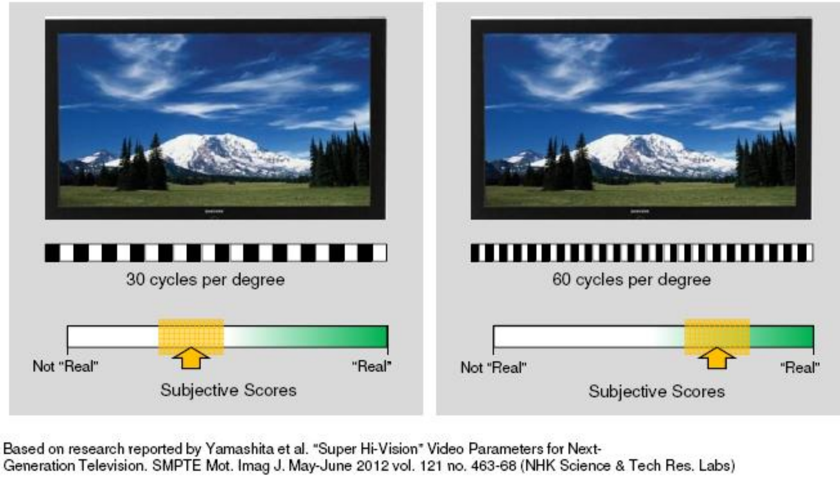

Figure 3 – "Normal" (20/20) Vision is Insufficient in Characterizing the Viewer's Subjective Video Experience [11, 18]

Studies also support the objective of a more *immersive* experience of a wide FOV (UDH-2). As shown in Figure 4, a wider FOV delivers the "being there" experience. To obtain this effect is a combination of a larger screen and close-up viewing, as in an IMAX theatre. And, of course, close-up viewing makes pixel density proportionally more important. According to ITU-R Report BT.2246: "UHDTV is … intended to provide viewers with … a wide field of view that virtually covers all of the human visual field."

The original UHD-2 recommendations by NHK are a careful balance of immersion and experience quality against the propensity for viewer discomfort as the brain tries to reconcile the imperfections of the artificial immersion environment.

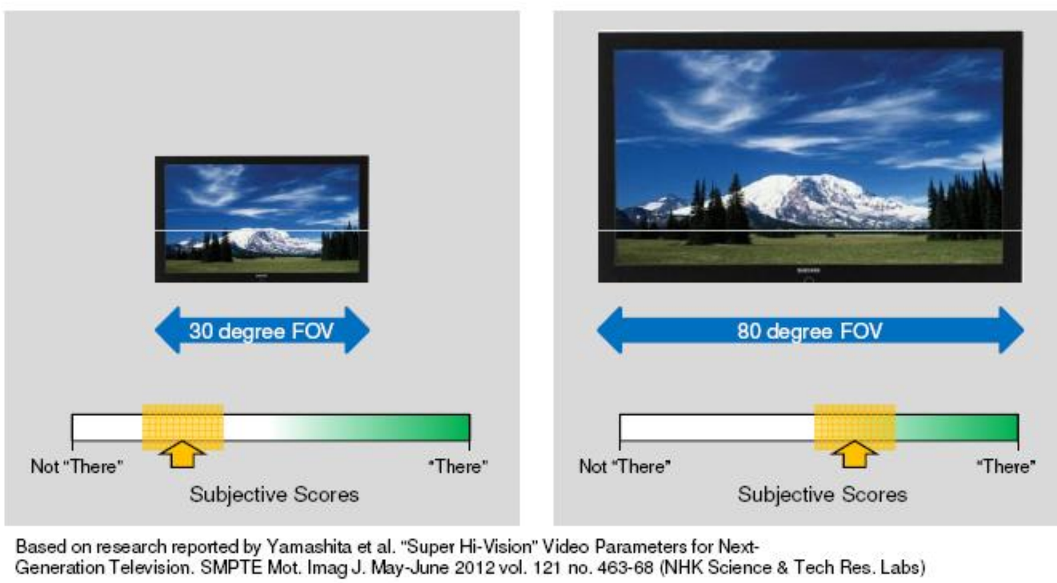

Figure 4 – Field-of-View and "Being There" Immersion [11, 18]

In addition to bit rates that will increase as pixel counts rise, a key component needed for quantifying screen sizes impacts for capacity management analysis is usage metrics. For UHD-2, the usage is constrained by its relationship to extremely large displays when compared to the inherent limitations of normal walls in a normal home. A 60" display is a five foot diagonal screen, and therefore a horizontal length of over 4 feet, while an 84" display (7 feet) has a horizontal length of over 6 feet! Though these display sizes (and projector systems) are suited to home theatres, they are not consistent with the living area viewing environment typical homes. They are very imposing companions in a normal living room. It is therefore deemed unlikely in our analysis that typical residential entertainment evolves in this way to take over the home.

In summary then, the long-term assumption that will be represented in subsequent capacity management analysis is that UHD-1 is a mass-market, scalable service for MSOs to introduce when the time is right. Data suggests this might be sooner than expected, with Credit Suisse projecting adoption by 135 broadcasters by 2017 [22]. We assume that UHD-2 is a service MSOs want to offer to the class of customers capable of enjoying it, but as an available IP program stream (or VOD title) only so as to avoid large spectrum penalties of broadcast.

Small Screens Pack a Big Punch

The HD 2.0/3.0 physics is not constrained to the primary screen. Consider Figure 5. Not only do larger primary screens suggest better spatial resolution, our secondary screens are now capable of high quality video such as HD. Tablets have changed the game for 2nd screen viewing, where a few hundred kilobits-per-second of low-quality video is no longer the norm for over-the-top (OTT) services. This changes how we quantify the impact of 2nd screen video devices receiving IP video around the home.

It is easy to see from Figure 5 that standard HD can be improved upon for reasonable viewing environments. For a 10" tablet, if the screen is about 17" away (airplane, back seat of a car), its spatial resolution can be perceptibly improved with a higher resolution format.

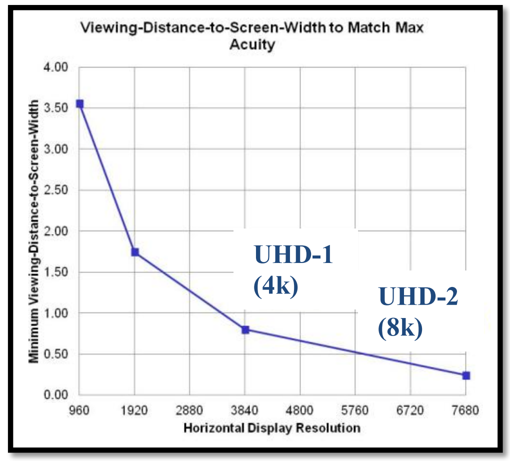

Figure 5 – Screen Size, Distance, and Resolution – Mobile Viewing

The case for UHD-1 based on the 10" tablet is difficult to make, but clearly screen sizes and portability have combined to change the paradigm of mediocre 2nd screen video.

The Need for Speed

The 30 Hz (interlaced), 50 Hz, and 60 Hz frame rates have origins in AC line rates, and thus are only loosely scientifically tied to video observation and testing. They simply were high enough to avoid the known issue of flicker. However, as spatial resolution has continued to improve, temporal resolution has not. Interlaced video itself is a nod to overcoming poor motion representation – exchanging spatial resolution for a higher rate of image repetition to better represent motion than a progressive scanning system of the same bandwidth.

Due to this legacy, HDTV today for progressive format is p60, or 60 fps. As displays become larger and of higher resolution and contrast, the challenges to effectively displaying motion increase because the edges to which movement is ascribed are sharper. Studies have shown that better scores for video QoE are given for higher frame rates for the type of video that intuitively would benefit most (high action, sports). Figure 6 shows such an example taken from ITU studies aimed at UHDTV recommendations.

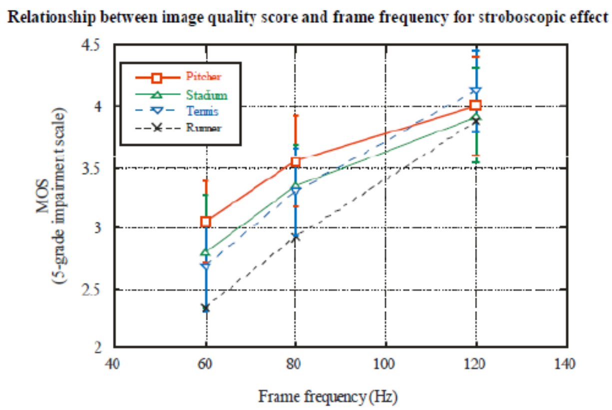

Figure 6 – Mean Opinion Score vs. Frame Rate and Video Type (Source: Report ITU-R BT.2246-1, 2012)

ITU recommendations for UHD include frame rates up to 120 fps. In addition to simply better representing high motion video, further ITU studies reveal that viewers are more susceptible to flicker for a wide Field-of-View (FOV). This suggests the need for higher frame rates for very large displays. The ITU judged 120 fps as necessary to minimize motion blur, stroboscopic effects, and perception of flicker for a wide FOV. Figure 7 from the same ITU report shows how the "flicker" frequency observable increases with a wider FOV.

Other studies have made similar conclusions about the perceptual benefits of increasing frame rates, including results as high as 300 fps [1].

We will consider 120 Hz as a component of the video evolution and quantify the effect. We will base use of this frame rate on the statistics of content type drawn from the most recent Motorola Media Barometer [24]. In this survey, 24% of viewing is sports, and another 24% is non-drama entertainment. We will consider these two categories as eligible for 120 fps. We thus assume that for UHD-1 viewing, eventually, 50% of it will be 120 fps. We will assume all UHD-2 viewing is 120 fps.

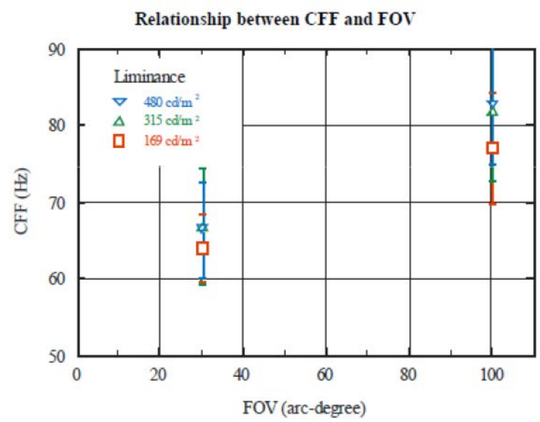

Figure 7 – Field of View vs. Critical Flicker Frequency (Source: Report ITU-R BT.2246-1, 2012)

ASSEMBLING LIABILITIES

Whereas 10-12 standard definition (SD) programs fit in a single 6 MHz QAM bandwidth, this number drops to 2-4 for today's HD, HD 1.0 – 720p and 1080i. Even at that, four is understood to be a nod to the trading off of video quality in favor of efficiency. Even worse, HD today represents a simulcast – programs delivered in HD are also transmitted in SD.

Moving beyond SD and HD1.0, UHD-1 (4K) works out to 4x the pixel count as 1080 HD, and UHD-2 (8K) works out to 16x the pixel count. These are major new Liabilities.

Relative to SD, we can summarize the pixel related bandwidth multipliers as:

1080i – 4x
1080p – 8x
UHD-1 (4K) – 32x
UHD-2 (8K) – 128x

Figure 8 captures the relationship amongst the formats. Note that the "Digital Cinema" 4K shown in green in Figure 8 is slightly wider than the format being discussed here (4096 vs. 3840), and thus the quotes around "UHD-1."

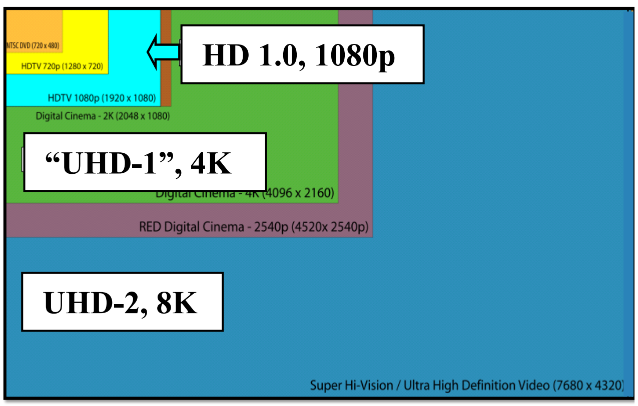

Figure 8 – Beyond HD 1.0: UHD-1 (4K) and UHD-2 (8K) [5]

The frame rate increase discussed is also an important Liability, scaling UHD-2 and, as described, about half of the UHD-1 programming.

There is one more Liability that we entertain. In UHD cases, there is the possibility of using higher bit depth (from 8-bit to 10-bit or 12-bit) color quantization, which also effects the bandwidth. In fact, the ITU Recommendation includes consideration of both 10-bit and 12-bit quantization depth. In the analysis, we apply 10-bit depth to UHD-1 only when high frame rate formats are used, and always for UHD-2.

All scale factors do not necessarily directly translate to bandwidth multipliers, but it is a useful first-order upper limit assumption. Our assembled *Liabilities* are therefore:

-1080p scaling from HD 1.0
-UHD-1 pixel scaling
-UHD-2 pixel scaling
-Frame Rate increase (some content)
-Quantization Depth increase (some content)

ASSEMBLING ASSETS

Encoding Efficiencies: H.264 & H.265

It has been 10 years since the Advanced Video Coding (AVC) [16, 18] international standard was completed in 2003. AVC – also known as H.264 and as MPEG-4 part 10 – and its equally successful predecessor, MPEG-2, are expected to continue to play an important role in the digital video economy for many more years, but they have been joined by the latest encoding standard -- High-Efficiency Video Coding (HEVC) [2, 3, 10, 14, 19, 20], or H.265.

The Final Draft International Standard (FDIS) for HEVC was ratified in January 2013. A deeper description of HEVC itself is given in [4]. Figure 9 captures the state of the set of core MPEG compression standards in the context of their lifecycle.

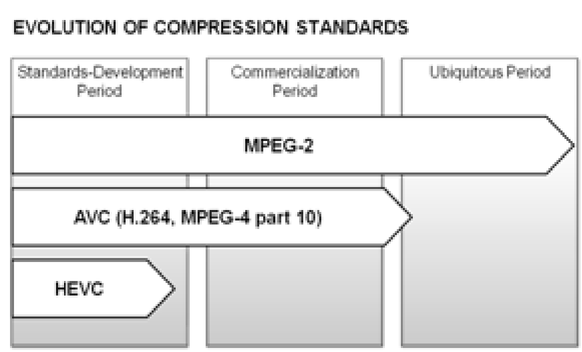

Figure 9 – State of Video Compression Standards [4]

H.265 roughly doubles the compression efficiency over its H.264 predecessor. H.264

itself doubled compression efficiency compared to MPEG-2. Thus, an 18 Mbps HD program using MPEG-2 would need only 4.5 Mbps using HEVC. Alternatively, UHD-1 can be delivered with HEVC at the same rate that HD over MPEG-2 is today. The savings represent key capacity management evolution Assets.

We will use 17 Mbps as our average data rate for a UHD-1 program at 60 fps using nominal bit depth (8 bits), based on recent internal studies of HEVC codec performance.

<u>New Spectrum Considerations</u>

Figure 10 illustrates the anticipated spectrum migration of the HFC architecture long-term. Because of many reasons outlined in [8] and [9], we foresee a phased approach to spectrum migration, consistent with the way operators incrementally deal with infrastructure changes in the context of dealing with legacy services and subscribers.

The end state of the spectrum migration is shown in the bottom illustration of Figure 10. This long-term end state (as an HFC-style architecture) maintains a level of asymmetry consistent with historically observed downstream/upstream traffic ratios.

Note that the objectives set out for DOCSIS 3.1 is to ensure at least 10 Gbps downstream and 1 Gbps upstream, and this phased evolution plan is meant to be aligned with that objective.

We evaluate spectrum evolution in our analysis of new capacity Assets, considering both the "Excess Bandwidth" case of 1.2 GHz and the extended bandwidth case of 1.6 GHz. We refer to a 1.2 GHz upgrade in subsequent tables as "Spectrum A" and 1.6 GHz as "Spectrum B."

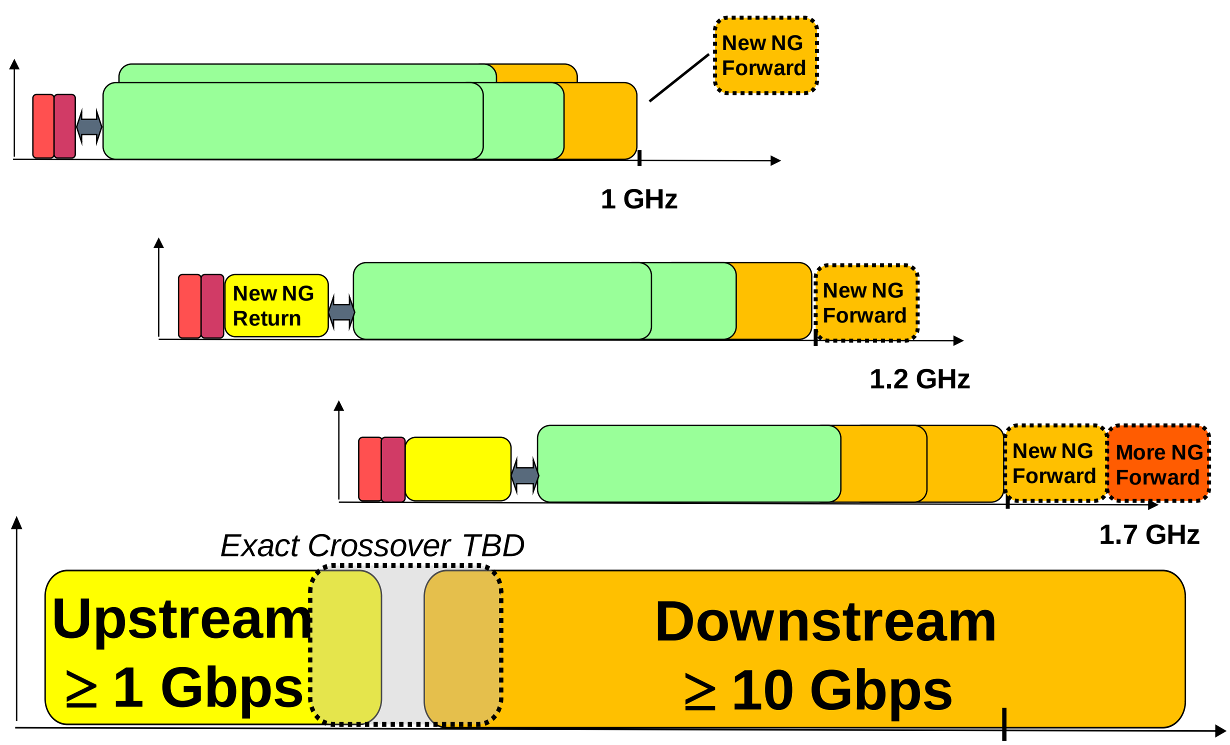

Figure 10 – Probable Evolution of the Cable Spectrum

THE dBALANCE SHEET – SIMPLE

In assembling our Liabilities, we quantified pixel count increases and accompanying video QoE parameters (frame rate and quantization depth) that translate into new bandwidth requirements. Going no further with media consumption bandwidth, these become our fundamental HFC capacity "Liabilities." What do we mean by going no further? We mean that we are excluding futuristic cases such as holograms or multi-dimensional displays. It seems reasonable to view these as beyond 15 years if they come to pass at all.

Also note that the analysis is limited to media consumption only, and not intended to address other potential services that may evolve to consume bandwidth. In our subsequent analysis, we *do* account for a separate Internet data service, but we do not attempt to quantify, for example, remote healthcare services, machine-to-machine activities, or other potential uses of the network that could affect bandwidth. We do not search for new "killer apps."

Assembling all of the Assets and Liabilities discussed above in Table 1 shows the Capacity Management Balance Sheet. At first glance, it appears threateningly out of balance from an operator perspective.

Table 1 – Capacity Management Balance Sheet, Simple Form

Assets	dB	Liabilities	dB
H.264	3.00	1080p60	3.00
H.265	3.00	UHD-1 (4k)	6.00
Spectrum A	2.43	UHD-2 (8k)	6.00
Total A	8.43	10-bit	0.97
Spectrum B	3.19	Frame Rate	3.00
Total B	9.19	**Total**	18.97

What does the table tell us exactly? By simply aggregating the columns, Table 1 suggests that if we take a snapshot of today's digital services that exist in the spectrum available (this example assumes a 750 MHz network – a value needed in order to determine new Spectrum Assets) and, fast forwarding to a future state where all are converted to UHD-2, High Frame Rate, enhanced bit depth, we'd have a (very) negative balance and thus not enough capacity to do so. It would be a great sign if the conclusion of such tabulation left us with a positive balance. But, since it does not, we ask: is this an evolution example that we really care about? In practice, the picture is not really this ugly – we just need to delve beyond this oversimplified arithmetic and analyze the problem a little deeper.

First, let's recognize that the more likely expectation, the long-term mass-market evolution assumption previously stated: UHD-1 (4K) takes hold as the significant wide penetration video service, with UHD-2 (8K) being an available format in a selective on-demand or VOD unicast fashion. We certainly do not envision a video evolution that takes all of the current programming line-up on the wire and converts it to UHD-2 to be put on the wire, which is what direct application of Balance Sheet arithmetic simulates.

Now consider Table 2. It quantifies the Balance Sheet outcome for various permutations of spectrum and service rather than this unrealistic one. Under the UHD-1 only scenario, as shown in Table 2, we identify a sliver of hope in the upper right-most cell of the table. In this case, assets just eclipse liabilities, the preference for capacity planning (clearly our version of a Balance Sheet would trouble actual Accountants!). However, we have yet to try and accommodate a portion of UHD-1 programs at higher video quality (VQ) such as frame rate or quantization depth. We can project that to do so would not be possible in any significant degree from a capacity perspective.

Table 2 – Achieving UHD-1 Balance

	Fixed Spectrum	Spectrum A	Spectrum B
UHD-1 Only	-3.00	-0.32	0.49
UHD-2 Only	-9.00	-6.32	-5.51
UHD-1 HQ	-6.97	-4.29	-3.48
UHD-2 HQ	-12.97	-10.29	-9.48

Diving deeper still, note that the current spectrum usage is actually a simulcast mix of SD and HD programs. The calculation in Table 3 converts these all to a single format. In other words, two of the same format of the same program is what the calculations in Table 2 show. However, the lower format simulcast represents a smaller percentage of the spectrum to begin with so the impact is limited. Significantly, however, as shown in Table 3, it does improve the outcome, as the assets now exceed the liabilities in the case of a 1.2 GHz spectrum (second yellow cell) expansion instead of only the 1.7 GHz case. Progress!

Table 3 – Correcting for Simulcast

	Fixed Spectrum	Spectrum A	Spectrum B
UHD-1 Only	-1.90	0.78	1.59
UHD-2 Only	-7.90	-5.22	-4.41
UHD-1 HQ	-5.86	-3.19	-2.38
UHD-2 HQ	-11.86	-9.19	-8.38

We can further estimate that any system that has at least 60 analogs today (750 MHz network) that is within 3.25 dB of balance will also be sufficient if the analog simulcast is removed and donated to the digital video pool. This additional set of successfully balanced options is circled in red on Table 3. Though there is little margin to be overly comfortable between Assets and Liabilities in virtually all cases, that we can emerge in positive territory is encouraging. It is not surprising that the extremely powerful use of analog reclamation shows significant Balance Sheet benefits. More Progress!

Encouraging as things are looking, and as satisfyingly simple as Balance Sheet tabulation can be, Table 3 still does not tell a complete evolutionary story. We do not anticipate that the approach for even HD 2.0 would be to implement every legacy digital program in the lineup today as a UHD-1 program, either broadcast or on SDV. It would be part of the broader IP transition – from a timing standpoint as well as from an efficiency of spectrum standpoint.

It is also insufficient to focus on final outcomes of service evolution only. The transition plan, which accommodates simulcast and interim architectural phases are critical for capacity management. A more comprehensive analysis approach will be required.

<u>The Intersection of Video Services and IP Traffic Growth</u>

The downstream rate of traffic growth has been chugging along at about 50% per year. The Balance Sheets approach of video service conversion does not account for this dynamic – it basically assumes that while video evolves the remaining spectrum allocation remains static.

Of course, historically distinct video and IP services are not staying that way. Indeed, it is one of the evolutions of video services as we know it – in the form of over-the-top (OTT) services – that is currently the engine of this 50% year-over-year Compound Annual Growth Rate (CAGR). In addition, MSOs are currently migrating their own video services to IP. Some have already begun the process. The IP Transformation is expected to take an extended period of time, as revenue-producing legacy services will not quickly be removed. As such, while tabulating Assets and Liabilities is very insightful, it is important for investment planning purposes to view an evolution timeline that quantifies all of the moving parts of video *and* data. The analysis approach is more complicated, but is still entirely tractable.

An example of the conceptual timeline approach was described in great detail in [4] and [8], and a sample transition analysis timeline from [4] is shown in Figure 11. We will summarize the process here as analysis later in this paper will use a similar methodology. Please refer to [4] for a complete description of this case and [11] for a walk-through description of how the analysis approach works.

As a brief context, Figure 11 charts the 50% CAGR number commonly used for downstream traffic growth (yellow trajectory) against various thresholds of video service evolution and architecture evolution through the year 2030. An inherent assumption of the 50% CAGR *in this version* of analysis is that the 50% CAGR includes the MSO IP Video transition – it is not separately accounted for. Another way to look at this is that the engine of growth driving 50% CAGR has been OTT video services (i.e. Netflix) and the shift from OTT penetration and growth to MSO introduced services is an invisible shift of eyeballs relative to CAGR.

The trajectory is broken up into four segments with three discontinuities that represent three service group splits – by half, by half, and then down to N+0 (approximately one-third at that point).

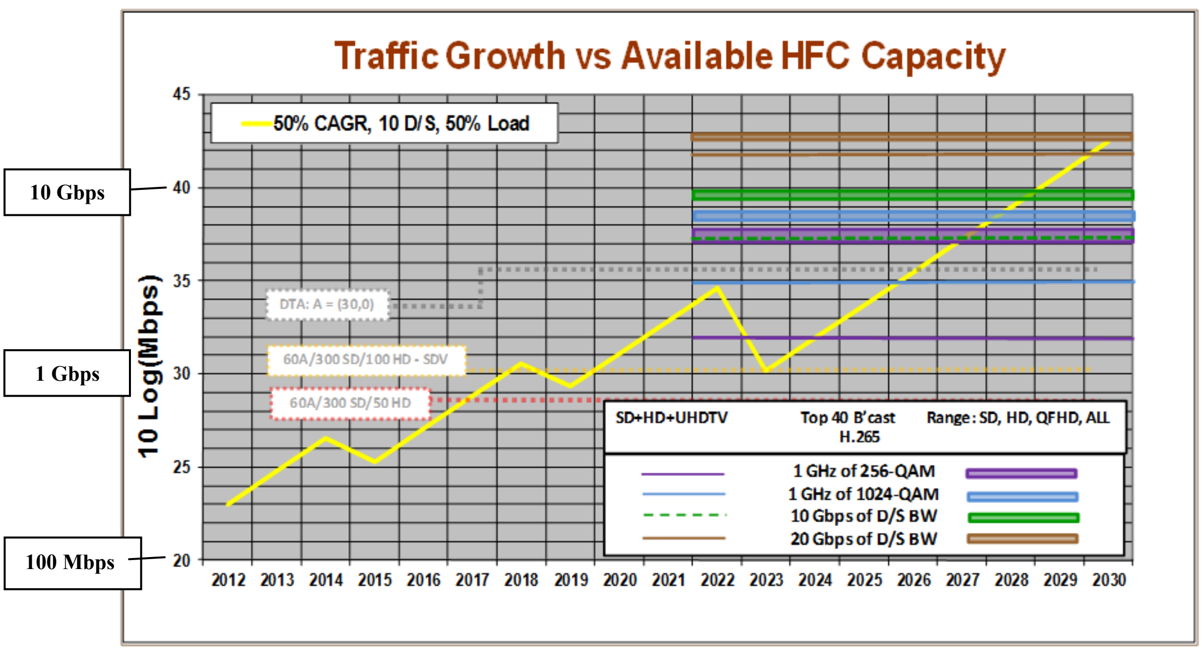

Figure 11 –HD 2.0 & 3.0, IP Traffic Growth and HFC Capacity Limitations [4]

Three labeled horizontal thresholds are shown on the left side with specific video service and implementation assumptions. For example, the red threshold bar is for a system with 60 analog video carriers, 300 SD, and 50 HD services in an 870 MHz system. If these services held steady, then the growth of IP traffic would breach the threshold in 2017, assuming one node split occurs in the intervening time period. The two horizontal thresholds above this represent the addition of HD programming (to 100) while also deploying SDV to manage the bandwidth growth, and the use of analog reclamation in two phases: 60 to 30 and then in 2017 the removal of all analog.

The remaining thresholds in the upper right are described in the legend at the bottom right. In all cases, there is an assumption of a static IP broadcast of the Top 40 channels, based on analysis and reasoning we shall describe in a subsequent section. Four cases of video format combinations are analyzed with different network architecture assumptions. See [4].

A key conclusion from [4] around Figure 11 was that the UHD-2 (8K) as a format with a broadcast component had major implications without significant new capacity exploitation. Of course, as previously stated, we do not envision UHD-2 as such a service. By contrast, for UHD-1 (4K) scenarios, even the least capable case (1 GHz of 256-QAM) had an extended lifespan. This bodes well for the ability of tools already available to manage through an aggressive combination of video service evolution and persistent CAGR of IP traffic. Figure 11 makes it abundantly clear that for a persistently high CAGR over a long period of time, CAGR eventually wins, a fact we revisit and reevaluate in the next section.

We will use a similar tool and analysis approach in the calculations ahead.

ASYMPTOTIC ALL-IP TRANSFORMATION

In the Figure 11 analysis, we approached the problem of CAGR and video services as largely orthogonal services sharing common spectrum to be managed. We accounted for the video services in the analysis approach in [4], after conversion to IP, as a spectrum block that was therefore unavailable to be used for CAGR associated with the HSD service growth.

In this paper, we instead consider the IP video and data services as a composite to reconcile capacity constraints and CAGR over a 15-year transition period. A core reason for this is based on the common assumption that the engine of 50% CAGR in the last several years has been video services – albeit over-the-top (OTT) video services. MSO managed IP Video services are on deck, and it is a central premise of the analysis approach taken here that streaming video has been and will continue to drive CAGR. This leads to a foundational premise of the analysis going forward:

With our knowledge of HD2.0 and HD3.0 visual acuity relationships and technology availability, and the assumption of video-driven CAGR, we can set some realistic boundaries on asymptotic capacity requirements for a given serving group aggregate. Armed with an asymptote, we can project long term CAGR from the media consumption standpoint that tapers to this projected asymptote over the transition duration.

Our approach, therefore, is an analysis aligned with the principles of [5], where theoretical limits of video representation were piled one on top of the other to develop a worst case media consumption asymptote per household. Here, we remove the theoretical aspects and consider more practical and technological aspects of HD 2.0/3.0 over a workable business planning horizon. The objective we execute on will now be to find these boundaries and develop a migration approach aligned with what we discover, using our Asset and Liability line items to guide a transition plan.

For the IP Video world, there are additional implications and valuable Assets to assess that are associated with managing streams instead of “channels,” which we now describe.

Trends Effecting IP Video

The nature of video traffic being delivered is changing. There are many variables in play, virtually all of which are driving towards increasing unicast, and perhaps acting as coal for the CAGR engine:

- More content choice
- Time-shifting
- Trick play expectations
- Network DVR (nDVR),
- Video capable IP device proliferation (tablets and smartphones)
- Shrinking service groups

One important result of these shifts is that gains typically afforded by multicast capability or bandwidth reclamation gains associated commonly with SDV architectures begin to erode.

The benefits (or not) of multicast bandwidth savings can be determined quantifiably. A unique tool has been developed that takes into account the traffic, device, and format variables, as well as the known viewing behaviors of service group aggregates based on the history of SDV and IPTV networks. It quantifies these behaviors in the calculation engine to predict bandwidth and channel usage requirements. The tool is freely available for use on the web at http://www.motorola.com/Multicast-Unicast-Calculator/.

We will use this tool to determine to IP Video bandwidth requirements for services as we know them today – DOCSIS 3.0 with SD and HD 1.0 services. And, because the calculation engine is agnostic to formats specifically in favor of defined bit rates, we can use it with HD 2.0 and HD 3.0 inputs along with projected assumptions about the relative viewing behaviors of future simulcast services to determine IP Video bandwidth requirements in the UHDTV era.

Lastly, we can adapt the outputs to determine DOCSIS 3.1 requirements, adjusting for the increased spectral efficiency D3.1 entails.

IP Video Traffic Multiplexing Efficiencies

Legacy architectures are based on simple traffic management techniques that allot and enforce an average of 3.75 Mbps per SD video stream in to fit 10 streams (at least) per QAM carrier.

The introduction of DOCSIS 3.0 adds channel bonding to the toolkit. The net effect of bonding coupled with more streams/QAM of H.264 or H.265 is the ability to use the law of large numbers, reducing average bandwidth. Many independent streams competing for much more pipe capacity result in a self-averaging effect [6]. Smoothing out peaks and valleys are handled inherently by statistics operating with a capped VBR scheme, and most likely with Adaptive Streaming to ensure QoE. An example multiplex of independent peak-to-average video waveforms used in simulations to quantify this effect is shown in Figure 12.

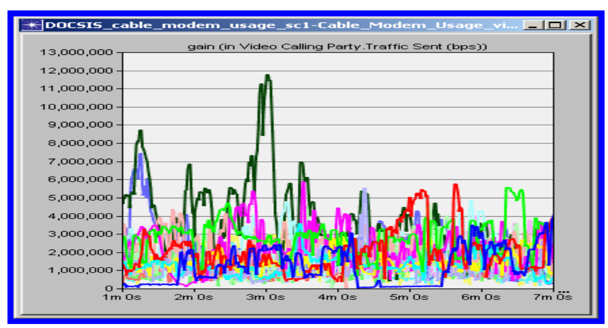

Figure 12 – Channel-Bonded VBR is a BW Efficiency Asset [6]

Based on simulations and observations, we use an 80% scaling as the bandwidth required for VBR-based channel bonded DOCSIS video in comparison to single carrier QAM transport.

Fiber Deep Migration

"Business as Usual" HFC migration has been shown to be well-suited to supporting a lifespan of at least a decade of legacy video evolution and aggressive IP data traffic growth [8]. The use of node splitting in the HFC architecture reaches its ultimate phase when the last active becomes a fiber optic node. This architecture goes by various names – Passive Coax, Fiber-to-the-Last-Active (FTLA), or N+0. Figure 13 illustrates this classic multi-phased operator migration strategy for segmenting a serving area.

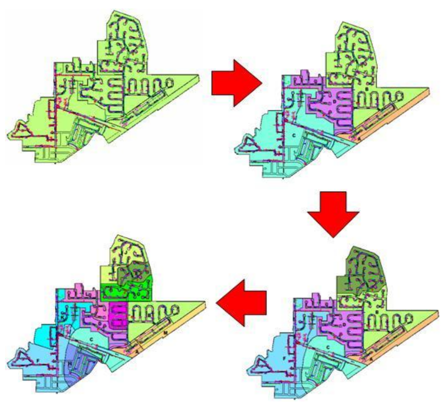

Figure 13 - Fiber Deep "Business as Usual" Migration is an Average BW/hp Asset

Regardless of the name, the architectural implications for N+0 have three core components:

1) Very small serving groups (40 hhp assumed)
2) The opportunity to exploit new coaxial bandwidth with no actives after the fiber optic node
3) A higher performing (higher SNR) HFC channel.

The latter two are both targeted by DOCSIS 3.1 – the first to enable the 10 Gbps worth of downstream capacity, the second to make use of the most bandwidth efficient modulation profiles (4096-QAM, possibly higher). This will offer the best opportunity to achieve 10 Gbps or more. Figure 14 compares today's 256-QAM format to the 4096-QAM format that clearly offers more bits per symbol, as envisioned in DOCSIS 3.1 for future bandwidth efficiency. They are shown at equivalent BERs of 1e-8 uncoded, which corresponds roughly to SNRs of 34 dB and 46 dB, respectively.

The N+0 architecture is essential for a growing forward spectrum to 1.6 GHz. Plant RF actives are unlikely to be stretched to 1.6 GHz, as they may be for the 1.2 GHz extension. N+0 will also leave operators within a stone's throw of FTTP.

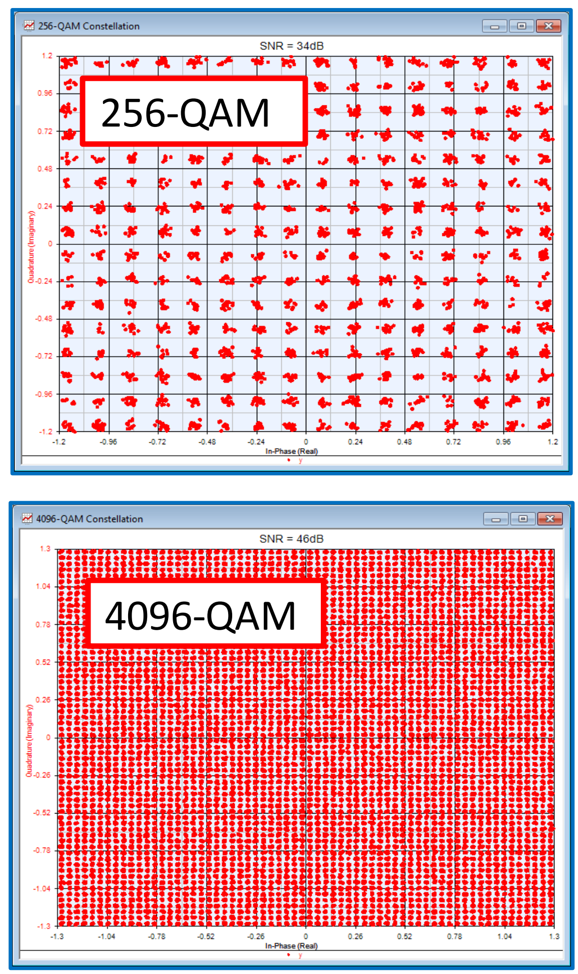

Figure 14 – DOCSIS 3.1 Modulation Formats as a BW Efficiency Asset

BEGIN WITH THE END IN MIND

An idea of where you want to be, to the best it can be known, is how to approach optimizing the path from A to B. In this case, "B" means an all-IP end state, and some projected mix of video services. Furthermore, we consider the end state as something accomplished over a 15-year transition period.

Importantly, in the all-IP case, we move away from channel-thinking and towards thinking in terms of streams. This is even more so the case considering DOCSIS 3.1. DOCSIS 3.1 will be a wideband OFDM system which literally removes the idea of channels in the conventional 6 MHz and 8 MHz sense. While DOCSIS 3.0 allows us to mathematically quantify the impact of wideband channels by allowing channel bonding, DOCSIS 3.1 remakes the physical layer QAM slots themselves in favor of OFDM subcarriers (also carrying QAM) able to be implemented as a single block over a very wide bandwidth.

What must be determined is "simply" how many streams and of what type to project as a 15-year assumption, from which the aggregate bits-per-second can be determined for the service group size envisioned. This is exactly what the modeling tool described above has an engine to calculate. However, the tool is for today's practical scenarios. This means a result based on DOCSIS 3.0 and input stimuli built around SD, 2nd screen viewing of varying bit rates, and HD 1.0. A snapshot of some of the parameters entered and used by the tool is shown in Table 4.

As noted previously, the stimuli used to drive the engine are merely names – any numbers can be put in that represent a scenario of interest, so long as the Mbps used in each category are aligned with the programming definition and usage defined for each.

Table 4 – IP Video Modeling Tool

STEP 1: DEFINE SERVICE GROUP SIZES, DOCSIS OVERHEADS, DOCSIS AND IP SERVICE PENETRATION

Service Group Size HHP:	125
Plus the following SG sizes:	62.5 and 31.25
DOCSIS Packet and Header Overhead:	5
DOCSIS Channel Bonding Overhead:	1
Overall DOCSIS Penetration:	70
Penetration Range of IPTV:	0.1 to 100 of DOCSIS Households
This Translates to an IPTV Penetration Range of:	0.07 to 70 of All Household Passed
Percentage of IPTV Customers Actively Viewing at Peak:	70% (50% is typical for cable systems as shown by historical SDV data)
Average Number of IP Devices per Home Active at Peak Viewing:	3.7

STEP 2: DEFINE THE TYPES OF VIDEO PROGRAMMING TO BE DELIVERED, THE REQUIRED DATA RATES AND THE VIEWING MIX BY CHANNEL TYPE

Required Data Rate by Program Type	Mb/sec
SD Bandwidth per Program MPEG4:	5
720p HD Bandwidth per Program MPEG4:	30
HD1080p/3DHD Bandwidth per Program MPEG4:	170
Medium Speed IPTV Bandwidth per Program MPEG4:	0
Low Speed IPTV Bandwidth per Program MPEG4:	0

We drive these inputs with the following video bit rate numbers and stream distribution assumptions, and also under an assumption of all HEVC at the end of the 15-year evolution:

For the 4K format, we will project that half are enhanced with 10-bit color depth and high frame rate (120 fps) with no additional encoding benefit assumed. More frames generally suggests less difference between frames and therefore potentially more coding gain, but the programming targeted is precisely the action-type video that is less likely to have that characteristic, or at least not to the same degree as a drama or news program. As we discussed previously, market analysis suggests this represents roughly 50% of viewing, so we account for this by having the enhanced UHD-1 as 50% of the total.

We project all of UHD-2 as 10-bit and higher frame rate, and assume this is a small percentage (10%) of pbh viewing associated with home theatre-type users. It is likely that the need for this service at all will be very serving-area sensitive (i.e. higher income suburban single-family home neighborhoods). Table 5 summarizes the values used as the viewing end-state parameters modeled.

Table 5 – End State (15-yr) Viewing Behaviors and Bit Rates

Format	Average Bit Rate	% Viewership
HD1.0	5 Mbps	20%
UHD-1	17 Mbps (50%)	70%
UHD-1 Enhanced	42.5 Mbps (50%)	
UHD-2	170 Mbps	10%

Two cases were examined for service group end states – operators who migrate to N+0 and others currently planning to be in the 100-150 hp "sweet spot" [25] which correlates to two more segmentations of the service group, whether virtual or actual fiber deeper. This leaves an "N+Small" HFC architecture in place, for example, such as an N+(1-3) cascade.

Other key assumptions include a 70% penetration (modest growth over the course of 15 years), a 70% peak-busy-hour (pbh) usage (aggressive), 1.5 streams per user per household, and the users per household governed by demographics associated with data extracted from recent 2010 census data, shown in Figure 15.

While we neglect mathematically household greater than 5 (all above 5 are treated as 5), we also neglect that about half of the households with children are those with children of an age unlikely to be independent viewers of multiple screens. There are other deeper weeds of demographic detail we could include such as this, but this model seems sufficient for better-than -ballpark estimates.

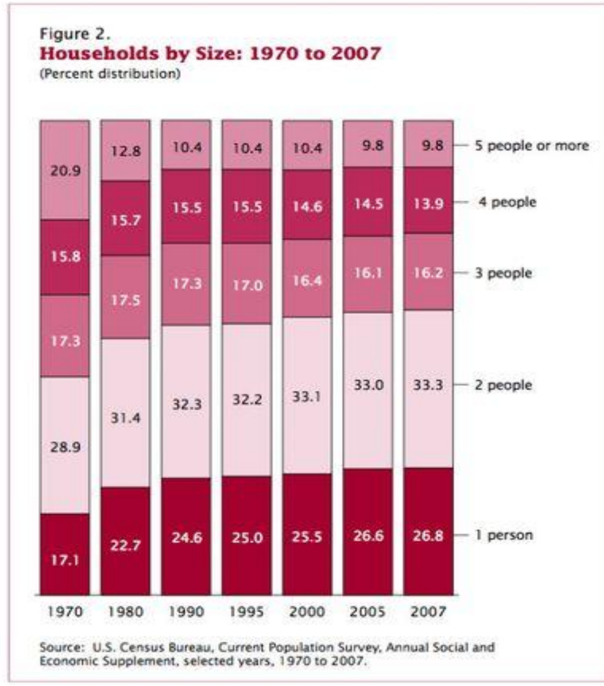

Figure 15 – Household Sizes to Govern Stream Counts [5]

Based on the above inputs for bit rates, viewership distribution, and service group size, we use a modified form of the model shown in Table 4 with HD 2.0 and HD 3.0 bit rates and usage metrics from Table 5. And, we recognize that the output is calculated in DOCSIS 3.0 channels, but which we can easily convert to Mbps or Gbps based on inputs for bit rate/channel of 256QAM and overhead losses. The two modeled cases and their results are shown in Tables 6 and 7.

Table 6 – End State (15-yr) Service Group Size of 125 HHP or "N + small"

Total DOCSIS Capacity Required to Deliver Managed IP Video Services

Total Bandwidth Requirements - Case 1

Service Group Size HHP:	125	Peak IP Video Devices / Home:	3.7 Active
Max Video Penetration:	70%	At Peak Penetration:	87.5 Homes
Percentage of Video Devices Active at Peak:	70%	At Peak Viewership:	227 Streams

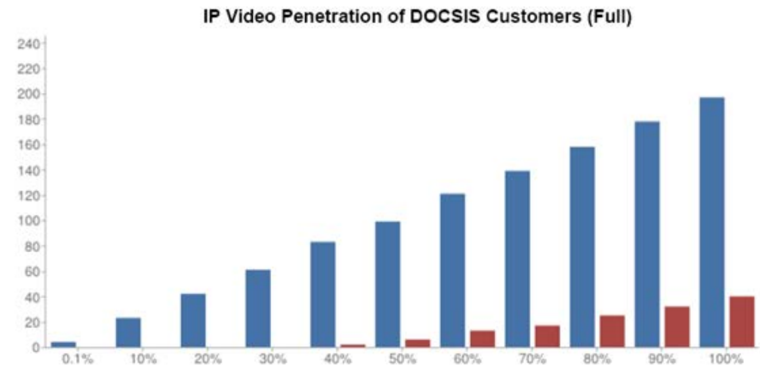

IP Video Penetration of DOCSIS Customers Charts

Total Downstream Channels @ 100% Unicast

Potential Channel Savings with Multicast

Table 7 – End State (15-yr) Service Groups Size of 40 HHP @ N+0

Total DOCSIS Capacity Required to Deliver Managed IP Video Services

Total Bandwidth Requirements - Case 1

Service Group Size HHP:	40	Peak IP Video Devices / Home:	3.7 Active
Max Video Penetration:	70%	At Peak Penetration:	28 Homes
Percentage of Video Devices Active at Peak:	70%	At Peak Viewership:	73 Streams

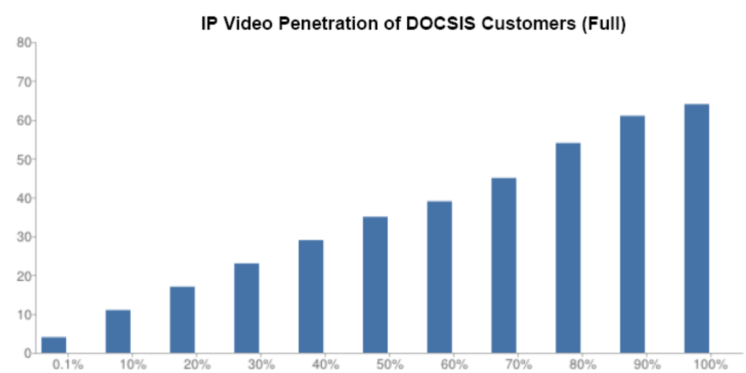

IP Video Penetration of DOCSIS Customers Charts

Total Downstream Channels @ 100% Unicast

Potential Channel Savings with Multicast

The analysis of the above two cases results in the following:

125 HHP (N+Small)

- 197 DOCSIS 3.0 Channels (~8.5 Gbps, overhead included)
- Potential savings of 40 channels available with multicast:
  - 157 D3.0 Channels (~6.7 Gbps)

40 HHP (N+0)

- 64 DOCSIS 3.0 Channels (~2.7 Gbps)
- No savings from multicast capability

Note from Table 7 that the serving group size combined with the unicast expectations and programming breadth has eliminated any multicast savings for the latter case.

Broadcast, Multicast, or Unicast?

One high level architecture result of this in-depth modeling of IP Video, and alluded to in the N+0 case above, is that because of the unicast trends, multicast gains may be limited and gradually erode. Because of this limited bandwidth benefit over time, and the complications brought about by multicast in the architecture (Wi-Fi, IP devices in the home, Adaptive Bit Rates), it may instead be simpler and nearly as effective to deploy a combined broadcast plus unicast architecture, since the vast majority of multicast gain is limited to the most popular programming. And, a broadcast component satisfies the "Superbowl" problem and even multi-channel major event scenarios (breaking major news story).

Analysis indicates that 80-90% of the multicast gain is obtained in the most popular 20 programs and fewer of course during "major event" scenario. This architectural concept is shown in Figure 16. We take a conservative approach in the analysis and will examine the case of 30 total broadcast programs, combined with the remaining all-unicast traffic as calculated by the model. We do not subtract any "broadcast" in the 125 HHP case where there is some available multicast gain, but will identify and compare the multicast example on the Capacity Management Timeline analysis. A detailed analysis of IP Video multicast and unicast architectures and implications is contained in [15].

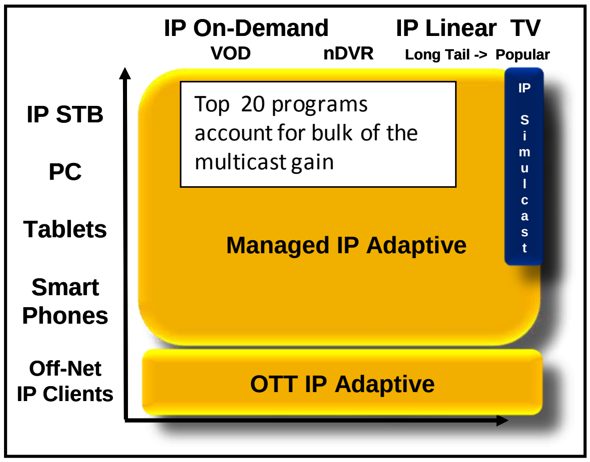

Figure 16 – Optimizing IP Video Delivery

Accounting for the IP Broadcast means allocating bandwidth for a simulcast of 30 IP channels in p60 HD and UHD-1 (4K), by assumption of desired broadcast service format mix in 15 years. For 4K content, we continue to assume half of the content benefits from enhanced quality (bit depth and high frame rate). Table 5 identifies the expected average bit rates this entails, and when it is aggregated over wideband IP channel it sums as shown in Table 8.

Table 8 – IPV Broadcast (30 Programs)

Format	Avg Bit Rate	Total
HD1.0	5	150.0
UHD-1	17	892.5
UHD-1 Enhanced	42.5	
Total	IPV Eff @ 80%	834.0

Under an assumption of *10 bps/Hz net* (payload) spectral efficiency we therefore allocate *85 MHz* for this broadcast spectrum. This efficiency is based on recent analysis done as part of the Channel Model ad-hoc in IEEE 802.3bn [26]. Analysis there has shown that over 90% of today's DOCSIS CPE report SNRs capable of 2048-QAM today, using current LDPC FEC technology to set QAM thresholds, as shown in Figure 17. 2048-QAM achieves a raw efficiency of 11 bits per symbol, and after allowing for overhead losses, a net efficiency of approximately 10 bps/Hz is expected.

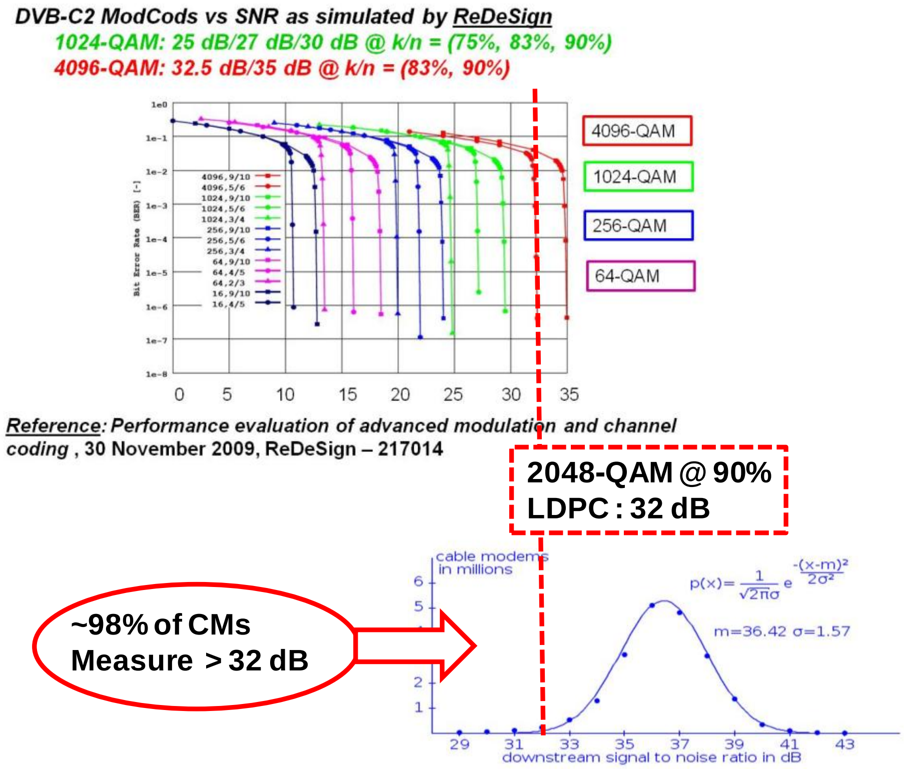

Figure 17 – Fielded CPE Report Much Higher Capacity Potential than Currently Implemented, with 90% Being 2048-QAM Capable [26]

Clearly, we will only march closer to capacity with FEC, and next generation CPE will be no worse in terms of sensitivity, fidelity, and implementation loss as today's – and should be significantly better.

Note that this implementation – broadcast + unicast – yields a total bits-per-second aggregate that represents a virtually non-blocking (Pb = .01%) unicast asymptote assuming our statistical basis of viewing.

<u>Let There be Data?</u>

Streaming video is the engine of CAGR growth in today's downstream, taking over for increased penetration of web browsing of sophisticated multimedia websites. This will remain the case with more bits required for HD 2.0 and the unicast driving trends previously described. However, in the analysis, we also set aside a complementary browsing experience. For this, we assume a

1 Gbps service with 1% concurrency (100:1 oversubscription) as a sufficient complement to the IP video service.

Given penetrations and serving group sizes, this works out to:

125 HHP: 875 Mbps (970 Mbps w/OH)
40 HHP**:** (N+0): 280 Mbps (310 Mbps)

We now have the components of our "End in Mind" requirements. We will analyze both by Balance Sheet and by the Capacity Management Timeline per the Figure 11 approach.

THE dBALANCE SHEET - REVISITED

We have now identified several new Assets and Liabilities in the prior sections. Let's add these to the Balance Sheet and assess their meaning. This is shown in Table 9 below.

Table 9 – Balance Sheet, New Entries

Assets	dB	Liabilities	dB
H.264	3.00	1080p60	3.00
H.265	3.00	UHD-1 (4k)	6.00
DOCSIS 3.1	1.76	UHD-2 (8k)	6.00
IPV Efficiency	0.97	10-bit	0.97
Split	3.00	Frame Rate	3.00
N+0	7.96	B'Cast	0.56
Spectrum A	2.43	**Total**	**18.97**
Total A	**22.12**	B'Cast A	0.33
Spectrum B	3.83	B'Cast B	0.24
Total B	**23.52**		

Here, "Broadcast X" represents the loss due to the choice to allocate 85 MHz of spectrum for 30 programs of IP Broadcast. Since this is an absolute set-aside, in this case an absolute network bandwidth was required. We again chose the 750 MHz downstream which incurs the most relative loss in the "Fixed Spectrum" case.

Repeating our matrix of possibilities, we can see in Table 10 that there are now *no deficit-only conditions*.

Table 10 – Capacity Balance, All Assets and Liabilities

	Fixed Spectrum	Spectrum A	Spectrum B
UHD-1 Only	10.13	12.79	14.28
UHD-2 Only	4.13	6.79	8.28
UHD-1 HQ	6.16	8.82	10.31
UHD-2 HQ	0.16	2.82	4.31

No deficit scenarios is an encouraging outcome. The massive benefit of service group splitting (nearly 11 dB total in the Asset column) is making a big difference. Of course, the use of all of the Assets listed in a capacity management calculation assumes a full transition to IP Video, where video streams and the associated traffic engineering of them allows the consideration of Asset parameters associated with service group sizing. Looked at another way, there is no bandwidth benefit to node splitting – to any serving group size – for broadcast spectrum. Segmentation of the network only has value in a switched architecture – whether that is classic legacy cable SDV or the IP Video case we are focusing on here.

Since IP Video *is* the plan for most MSOs long-term, *and* it is the end-game assumption we are working with here, the conclusions to draw from Table 10 should indeed be viewed as positive indicators for the long-term outlook. Table 10 suggests that there should be no scenarios where N+0 is an insufficient end-state solution. However, removing the N+0 (4.96 dB – replace it with a normal split) and perhaps some of the spectrum expansions (2-4 dB) would result in some negative balances in Table 10, suggesting capacity limitation for some of the 125 HHP scenarios.

Fully examining Tables 9 and 10, we find that this is so only for the cases where conversion to UHD-2 is the service assumption objective. As has been discussed, UHD-2 is viewed for this analysis as an available unicast streaming format to a small anticipated percentage of subscribers, and for which we have traffic engineered unicast IPV to support it based on Tables 6 and 7.

Encouraged by the long-term, the complete task then involves surviving the transition period of this capacity management challenge. Again, old services do not immediately die when new services are introduced. Likewise, new services rarely are introduced through massive service change all at once. However, allocating resources for them generally comes at service introduction. An IP approach helps to manage this resource allocation proportional to the penetration.

With the transition as the key challenge, we will walk through a timeline of potential service evolution-architectural migration scenario aligned with reasonable timings of each. First, however, we put our end state IP Video calculations to use and discuss their context and role in analyzing the transition process.

CAPACITY MANAGEMENT TIMELINE

With the above accounting for IP Video downstream bits and services, and our Asset and Liability line items to play with, we have all of the information we need to evaluate and develop a comprehensive Capacity Management Timeline for evolution planning. An important modification to the Figure 11 approach is used to bound problem by considering CAGR from a perspective other than blind allegiance to persistently aggressive growth.

Asymptotes of Behavior and Biology

There are *three key principles* to this perspective and implementation on a Capacity Management Timeline:

1) Recognition that detailed residential demographics are available and are more useful as metrics for an IP streaming world and multiple IP devices per home video connectivity than homes passed.

2) Recognition that humans have a limited ability to multi-task, in particular with video. While secondary screens as simultaneous and background content playing during a primary viewing experience may be common, humans have a limited ability to focus on multiple things at once with comprehension. Because of this, we cap simultaneous streams per home at $\leq$ 2 per individual (1.5 was used).

A counter argument to (2) is the “pub-style” home environment, with TVs just on 24/7 in rooms throughout a home, relatively independent of occupancy. The analysis bets against this, with “green” objectives perhaps a factor in this style of viewership evolution.

3) From the standpoint of *purely media consumption driven bandwidth*, or, alternatively aggregate bandwidth strongly dominated by media consumption, the IP Video tool output in the prior section coupled with the HSD service assumption calculation represents a *projected bandwidth growth asymptote*.

It is of course risky to suggest that bandwidth may stop growing, lest future work can look back and snicker about the naiveté of such an estimate – similar to early predications about the necessary memory requirements for a PC. Prior analysis [4, 5] considers assumptions otherwise, so consider this one in a series of analysis to weigh the possibilities and make a judgment. Of course, observation of trends over the course of time allow the industry to update the projections and react (if necessary) accordingly.

A logical reasoning besides visual acuity arguments for consumption bandwidth asymptotes can be built around the historical basis of much of the 50% CAGR number itself. From a media consumption standpoint, we can recognize that the speeds delivered from the Internet's mass-scaling outset were in large part associated with chasing the next required bit rate for increasingly higher level human media experiences:

1) Alphanumeric characters
2) Voice
3) Images (pictures)
4) Music
5) Low speed video
6) SD Video
7) HD 1.0

And now, per this analysis: HD 2.0 or HD 3.0 video. Based on our prior discussion on hyperacuity and studies such as [11], where it is foreseen that UHD delivers "retinal" image quality, and recognizing that normal households have wall sizes that limit the reasonable size of displays, there are several reasons to anticipate that *there is little practical benefit to services beyond the UHD format objective*s quantified in Table 8. In other words, our media consumption rate chase comes to an end (short of the holograms, etc.).

This then logically leads to our "asymptotic" limits for the service groups as shown in Table 11. Note that the IPV broadcast of 30 programs will be accounted for in a different manner on our Capacity Management Timeline. As we shall see, the broadcast spectrum offset plays a part in setting the threshold of available capacity for IP growth.

Table 11 – Projected Growth Asymptotes

	125 HHP	40 HHP
	(Mbps)	(Mbps)
Unicast Only	9470	3010
Multicast	7670	3010

Under these assumptions, it is insightful to recognize that the highest aggregate in Table 11 (125 HHP, Unicast only) is less than 10 Gbps, meaning it is lower than the objective called for as capacity for the DOCSIS 3.1 downstream – 10 Gbps or greater.

From these calculations, we can determine a 15-yr growth rate that ends up at 9470 Mbps/125 HHP from a starting point of 12 DOCSIS carriers (assumption) to a 500 hp serving group. The *resulting average CAGR for 15 years works out to about 33%.*

Over the 15 years, we have broken the growth rates into 3-year segments that could represent a possible play out of the entire 15-year transition. The incremental CAGRs used (average to 32.8%) are as follows:

Years: 1-3: 40%
Years: 4-6: 30%
Years: 7-9: 40%
Years: 10-12: 35%
Years: 13-15 (Complete Transition): 20%

The pattern recognizes that current CAGRs may settle after successive years of very rapid OTT growth. A new engine may be cable managed IP Video services penetrating alongside current OTT services, which will begin slowly and be many years in becoming highly penetrated, but will keep growth compounding steadily.

In the 7-9 year period, there will be some scale of 4K television sets (2015 being the threshold year) and a reason to move services in the direction of HD 2.0, increasing CAGR in the process before once again settling at the end of the cycle. This is when the final phase of the UHD service mix has been deployed, consumer usage patterns settled, and no clear growth engine for continued bandwidth expansion emerges for media consumption, at least as we know it and understand it today for business planning purposes.

Using the above CAGR segments and the boundary conditions calculated previously, the growth trajectories described and the asymptotes calculated can be visualized on the Capacity Management Timeline as the projection shown in Figure 18 for the two serving group sizes discussed. We have added two years at the end to allow us to envision a full tapering of CAGR. We will discuss the other Figure 18 markers and labels in the next section.

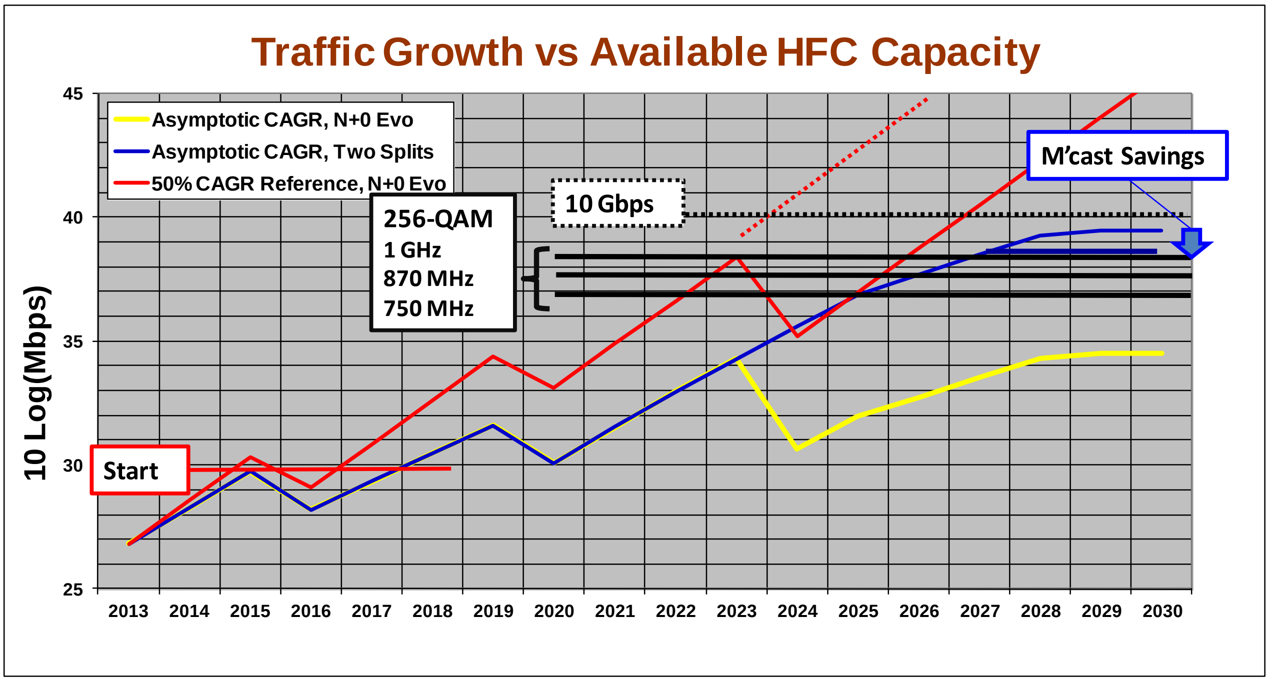

Figure 18 – Staggered CAGR Tapering Towards a Deterministic Video Services Evolution

As previously highlighted, it is very significant to point out that the assumptions of the CAGR rollercoaster ride and ultimate settling aligned with "retinal" video quality yields a sub-10 Gbps capacity aggregate. As Figure 18 shows, it does not leave very much margin beneath 10 Gbps for variations in assumptions, but then again DOCSIS 3.1's objective is a *minimum* of 10 Gbps.

Additionally, we have plotted the case that includes potential multicast savings – a small relative offset on Figure 18 as a decibel scale, but perhaps one that becomes important savings given the small margin between unicast and the 10 Gbps threshold.

Note that in prior work [5], as discussed, we estimated broader *theoretical* asymptotic boundaries of media consumption. The two key factors considered in [5] that are not addressed in this more practical perspective are consideration of the entire field of view of human visual system – because clearly sending more than this would not be sensible – and frame rates up to 300 fps based on some advanced research into how high the frame rate can go and yield an observable

VQ difference. Please see [5] for more details.

In this paper, we are taking the approach that such theoretical boundaries are unlikely to come into play in a time frame meaningful to be planning for them. A key conclusion even in that case was that, for an N+0 architecture, the aggregate *theoretical* capacity requirements could still be met, albeit with more aggressive RF distribution evolution [7]. This again is encouraging regarding our N+0 capacity fulfillment expectations.

Lastly, we re-emphasize that we are focusing on *media consumption-based service only as we understand them today*, and not trying to account for yet-to-be-determined bandwidth hungry applications, or even of the volume display or hologram sort of media consumption – considered outside the window of interest to compare.

AN EXAMPLE, PHASED MANAGED EVOLUTION OF SERVICES

Refer again to Figure 18. Several thresholds are drawn horizontally across the traffic trajectories shown.

On the left hand side ("Start"), the assumption used to guide the timeline is that we have fully utilized spectrum today. There is very little room to add new DOCSIS channels to accommodate growth. This is generally where MSOs are today in North America. We orchestrate the analysis in a way that is mostly agnostic to whether the network is 750 MHz or 870 MHz of downstream bandwidth. We only assume that in both cases, the spectrum is full. They will be different, obviously, in the types and amounts of services they offer. As a simple example, the 870 MHz network may carry 50 more broadcast HD channels.

The assumption threshold at "Start" is that up to 24 downstream DOCSIS channels can be squeezed out through some combination of tools – be that more SDV or removing some analog programming. In so doing, combined with a node segmentation plan over the next three years, the most aggressive CAGR situation stays below threshold for the duration, leaving a relatively short time window to execute on additional bandwidth recovery mechanisms. With the defined CAGR slow down (blue), this is extended by only about a year.

The upper right thresholds of Figure 18 are insightful. There are three traditional HFC downstream spectrum definitions, and the capacity associated with each if they were completely full of 256-QAM. These work out to, for the bit rate on the wire, as:

750 MHz – 116 slots; ~5.0 Gbps
870 MHz – 136 slots; ~5.8 Gbps
1 GHz – 158 slots; ~6.8 Gbps

Of course, these are not available capacities (yet) as they are consumed with legacy services. But, they offer immediate insight into the ability to architect for sufficient long-term capacity by simply comparing them to the trajectories in blue (125 HHP) and yellow (40 HHP).

The final threshold is the DOCSIS 3.1 objective of 10 Gbps. With DOCSIS 3.1 defining advanced modulation profiles and extended spectrum, there are multiple combinations of spectrum allocation and M-QAM order to achieve 10 Gbps.

+5 Years

As pointed out previously, while the end state of IP transformation may be encouraging, getting over the "simulcast bubble" is a complex capacity management challenge. Our groups of 256-QAM thresholds in Figure 18 look encouraging, but nearly all of that capacity is already spoken for. We have stitched together a possible scenario that looks at two 5-year snapshots of service and architecture evolution that encompasses current trends in service, HD 2.0 and HD 3.0 projections, and line items from the Assets toolkit.

Consider Figure 19. A new purple threshold range has been added for analog reclamation savings, with an assumption that 100% of analogs (60) will be removed at the end of five years. Most MSOs expect to remove all analog services, though fully doing so in 5 years may be aggressive for some. The analysis is easily adjustable for some small subset of analogs it may be desired to keep for traditional cable TV customers. However, it is readily apparent what the capacity bang-for-the-buck is in this most efficient of spectrum reclamation actions that can be taken.

The range associated with the analog reclamation rectangle is associated with the level of DOCSIS 3.1 migration – from none to an average 25% additional efficiency. This could be full migration enabled up to 1024-QAM, or partial with an average better than 1024-QAM on the D3.1 services.

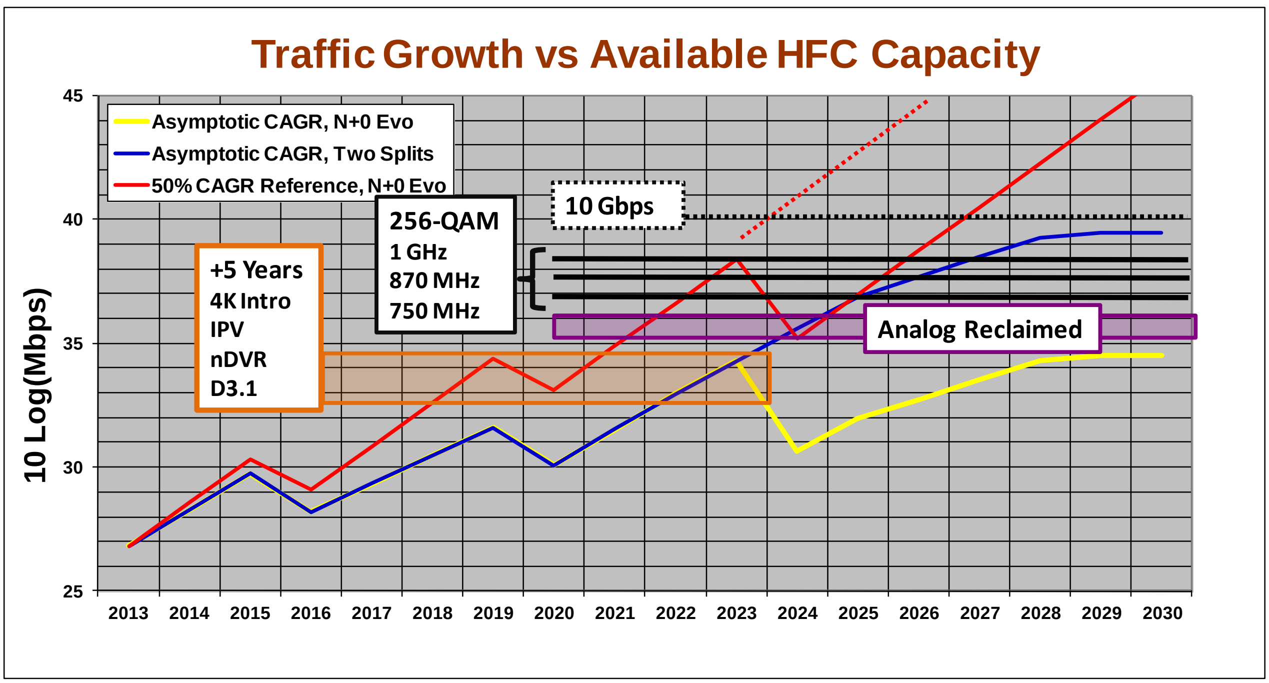

Figure 19 – Capacity Management Timeline – 5-year Snapshot

The orange rectangle represents what capacity remains for IP growth when the analog reclamation assumption is made, but with set-aside spectrum for legacy digital services and new video service evolution. The set of service and architecture evolutions assumed to have taken place by year 5 are:

Five-Year Snapshot

-Introduce 4K content into the VOD service offering (HEVC)
-Mix of VOD usage shift: 70/30 HD/SD to 30/60/10 UHD-1/HD/SD (no change in usage concurrency at pbh of 10%)
-Broadcast 10 programs in 4K (HEVC)
-IP Video Simulcast (25% penetrated, D3.0)

The same IP Video modeling tool previously extended for HD 2.0/3.0 services was used in its more common role of determining the number of DOCSIS 3.0 slots needed to support the legacy video mix, with the analysis resulting in 19 DOCSIS 3.0 slots required (and again not very much to be gained by multicast). At this phase of in the timeline, 250 HHP is the state of the serving group size. The results are shown in Table 12.

Table 12 – IPV BW Rolled Out over D3.0

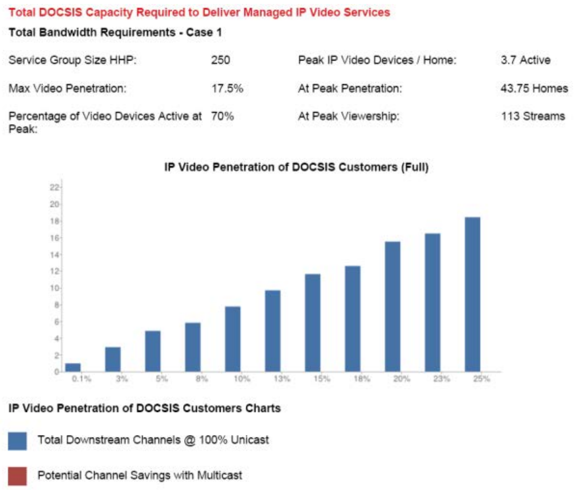

Creating the range of variation (the rectangle) are two considerations not deterministically assumed:

1) DOCSIS 3.1 penetrated in a meaningful way, or not at all. When so, the calculation assumption is enablement of 1024-QAM for bandwidth efficiency.

2) Rolling out of network DVR services, or not. Recorded content today in the US represents about 1/3 of the total content viewed [24]. Guaranteed unicast concurrency of video services today represented by VOD will increase with nDVR to this value plus VOD, worst case. We assume at +5 years that it has risen to 20% of viewing in addition to VOD. The same viewing format mix is assumed. This represents an aggressive assumption given the limited UHD-1 content available to record at this stage of the service evolution.

Observing Figure 19, what can we conclude? First, to make even reasonable room, a complete analog reclamation was assumed, and the 5-year plan may be aggressive for that assumption. However, it is likely that partial analog reclamation at least can be assumed.

Despite the resources freed up, the simulcast bubble that keeps legacy QAM in place while adding new bandwidth-draining services may leave the network vulnerable under persistently aggressive 50% CAGR at the end of the decade (2018). Implementing a second split by decade's end, along with gains of DOCSIS 3.1 and removal of nDVR from the mix extends this by about three years (2021).

Alternatively, a settling of the aggressive 50% CAGR, as shown in the blue trajectory, does not threaten the entire range of "thresholds" in the orange box for the 125 HHP serving group sized over the next nine years. That is a comforting window of time, and emphasizes why it is also important to keep a continual eye on CAGR trends. Again, the introduction of DOCSIS 3.0 IP video does not create new eyeballs in the service group; it (mostly) shifts them. Accounting for the spectrum allotted to IP Video channels (19) *and* assuming full-speed ahead with 50% CAGR may be an unrealistic exercise in double counting. A settled CAGR coupled with the pre-allocated IP Video traffic-engineered spectrum may better describe the dynamics.

Overall, the subtle lag in the CAGR per the prior assumptions (6-yr CAGR becomes 35%) pushed the "Year 5" service mix out to at least 8 years before the modified trajectory breaches the capacity barrier. Without the introduction of nDVR, this extends to 10 years. And, of course, with a less aggressive nDVR assumption – either the viewing mix skewed more towards traditional HD or more modest penetration – something in-between would result.

The 5-year snapshot also points out how subscribers might first get a taste of the new UHD-1 service in VOD and limited broadcast. By doing so, MSOs can observe the reaction and interest, and determine possible ARPU avenues in scaling the hot new format accordingly. Capacity management at this stage involves primarily limiting early programming, expecting segmentation, and removing analog carriers. In this case, we have left two levers – the Asset of DOCSIS 3.1 and the Liability of nDVR – as examples of how strategic decisions reflecting the pace of service and technology infrastructure investment can impact network lifespan.

<u>+ 10 Years</u>

Now consider Figure 20.

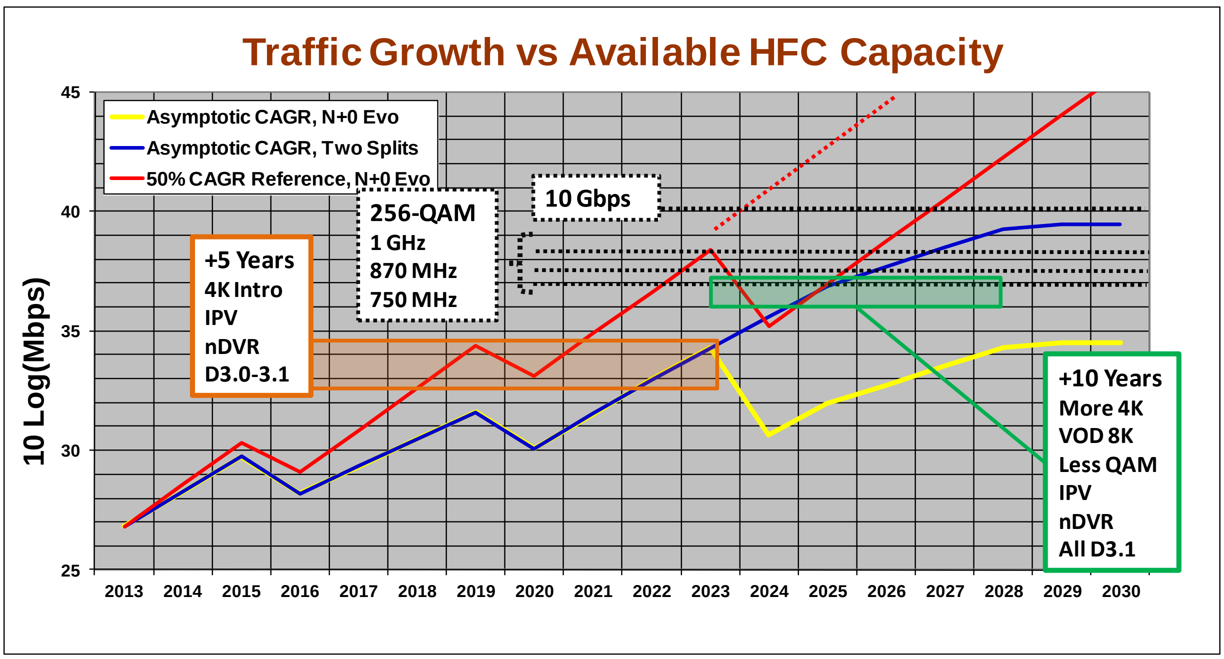

Figure 20 – Capacity Management Timeline – 5 and 10 year Snapshots

Here, a green rectangle represents a snapshot at +10 years, where the assumptions of the evolution state are as follows:

Ten-Year Snapshot

-VOD Mix 70% UHD-1, 25% HD, and 5% UHD-2 (8K home theatre, always HQ)
-VOD all IP/DOCSIS 3.1 (IP Unicast)
-Broadcast 30 programs in 4K (HEVC); 50% HQ (high frame rate and 10-bit), moved to IP
-IP Video is 50% penetrated
-All of DOCSIS 3.0 to 3.1
-DOCSIS 3.1 assumption 2048-QAM avg
-Total recorded viewing (nDVR traffic engineering) @ 30%
-50 legacy SD and HD channels broadcast

Most of the increase in available capacity comes from the massive reduction of legacy QAM content, making room for what now is a sizable broadcast of UHD-1 4K programming. The 50 programs of broadcast are along the lines of maintaining the basic service subscriber's offering today, just translated to the era of "you-can-no-longer-avoid-digital." There are also savings in D3.0 to D3.1 bandwidth efficiencies that are meaningful, though not as large.

Note that the entire programming offering of course still exists in SD and HD, it has just been migrated to IP, as has the 4K HD, which presumably has passed its consumer interest test at Year 5 as a service to scale to mass consumption.

The capacity gap between the orange and green rectangles identifies new lifespan to be engineered should the potential breach in capacity threshold that projects in Figure 19 come to fruition. Of course, the evolutionary steps are not discrete as pictured, and steps towards the ten year rectangle of capacity can be made that basically bridge the gap in the figure with a continuous threshold moving "Northeast" on Figure 20.

These bandwidth reclamation measures for the 125 HHP case, however, run out of steam by 2026, and there is no path under these assumptions that suggest two service splits suffice under either the continued aggressive CAGR of 50%, or the settled CAGR case which asymptotes just below 10 Gbps.

However, it is also notable that a prior-to-2024 evolution to N+0 meets the requirements of settled CAGR capacity growth, though not a persistently aggressive (red trajectory) 50% CAGR. Of course, this is really simply testament to the axiom that without infinite bandwidth, over time, CAGR always wins – it is just a matter of the time scale chosen.

In summary, for this 10-yr snapshot and assumptions, we have shown that an N+0 evolution meets the capacity needs of long term HD 2.0/3.0 evolution coupled with an HSD service, and meets the transitional needs of the simulcast period prior to getting to the asymptotic point. We have more deeply penetrated new services and technology by Year 10, while reducing legacy service offerings in a managed way that stays within available capacity with some level of certainty in the N+0 case. There is a 10+ year time window to evaluate whether a 125 HHP "sweet spot" is sweet enough to emerge as sufficient for the years that follow.

Network Nirvana

We calculated our 15-year end state IP transformation previously. There, legacy-free IP video bandwidth requirements under HD 2.0/3.0 evolution were determined using the modeling tool extrapolated to the UHD generation of video services.

As we did in the transitional 5-yr and 10-yr snapshots, we now account for the set aside spectrum at Year 15, in this case for 30 broadcast IP programs in UHD-1 and 1080p60 HD. Again, our UHD-1 assumptions are that approximately half of the programming would benefit from the enhanced video quality. UHD-2 is available as a unicast IP Video stream to the home theatre crowd. Recall, we calculated that the spectrum required for the IP broadcast would be 85 MHz.

Lastly, to truly account for all long-term network expectations, we must account for the expanded upstream, which may top out at 200 MHz. It is anticipated that an 85 MHz mid-split will be the first phase of upstream evolution, which has an impact on the downstream as well, but which is relatively modest. A life expectancy of at least ten years is foreseen at upstream growth rates [8], [9] under two splits and 85 MHz.

However, we will assume the more aggressive case for the "high split" comes to pass by Year 15. It is consistent with the DOCSIS 3.1 objective of enabling 1 Gbps of capacity upstream. Obviously, it also represents the most threatening case for managing downstream growth, since significant downstream spectrum is now sacrificed. In fact, the 1.2 GHz forward band was first conceived as a way to offset the loss of downstream spectrum as the upstream expanded [7].

For this Year 15 case, refer to Figure 21.

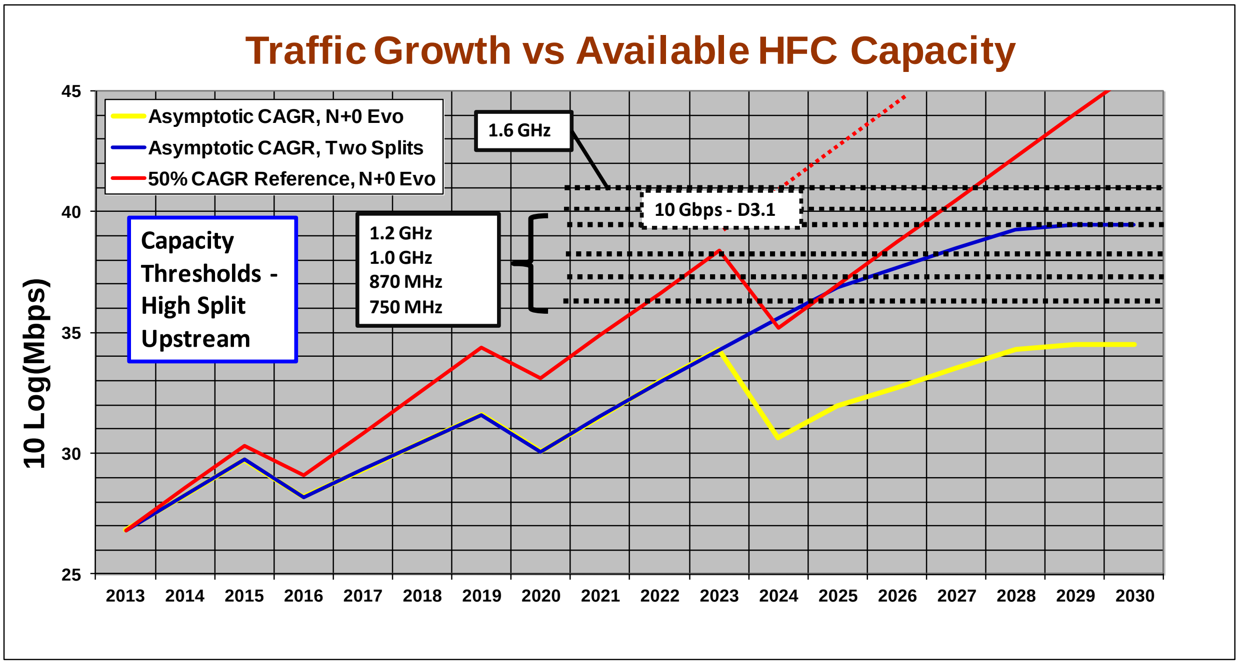

Figure 21 – Capacity Management Timeline
15 Yrs, IP Broadcast/Unicast Architecture, Upstream High Split

Figure 21 identifies the new capacity thresholds under the assumptions of the unavailable 85 MHz broadcast spectrum and spectrum donated to the upstream for several forward spectrum scenarios. The figure assumed that the forward path would begin at 250 MHz and achieves 2048-QAM.

As might be expected from all of the clues we have accumulated in our Balance Sheet tables and calculations along the way, an evolution to N+0 is sufficient from a capacity perspective in the end state of all-IP transformation, with asymptotic consumption behavior, even for the most constraining of forward bandwidth scenarios, although the gap closes noticeably with decreasing available downstream.

As important, it is *not* the case that capacity is always sufficient when the segmentation is limited to two splits. *Without bandwidth extension, capacity is insufficient*. Under the asymptotic growth scenario and an extension to 1.2 GHz, the capacity threshold and consumption asymptote are basically identical. The many variables in play could swing that scenario either way. For example, we might assume DOCSIS 3.1 makes it to 4096-QAM for all or enough to move the bar, or will we may squeeze a little more efficiency out of HEVC than foreseen.

However, the large impact variables involve assumptions of the average CAGR over long periods of time, and percentage viewership metrics for the most bandwidth-consumptive HD 2.0 and HD 3.0 streams. The accuracy of these assumptions is more likely to determine the sufficiency of defined thresholds, and especially so for borderline cases.

Perhaps most importantly, like any capacity management analysis, the work is a living document, with periodic updates associated with trends observed and technology shifts necessary to adapt the path forward.

SUMMARY

In this paper, we executed a long-term capacity management analysis, with permutations of scenarios of current and future services and architectures. While such long time windows are sensitive to assumptions, it is important to understand that all of the possibilities are quantifiable in straightforward fashion. The analysis undertaken here quantified each individual evolution variable in terms of its role as Liability or Asset, considering the extended time period of an HD 2.0/3.0 evolution and IP Transformation, and described the intricacies of deploying these Balance Sheet line items in a Capacity Management Timeline analysis.

From our perch today, the end state appears to be an attractive one in terms of available capacity for projected services over evolved HFC. We also examined a phased example transition to visualize the complex capacity balance and timing involved in crafting effective migration strategies. We reemphasize that, as insightful as the results herein may be, long-term capacity analysis is very much a living exercise and our team expects to continue to update our perspective periodically. Nonetheless, being able to comprehensively understand and methodically quantify the problem is essential to properly engage in effective scenario planning at every stage of the exercise, and to enable optimization of solution paths suited to an operator's particular circumstances. We hope this paper helps the industry in exactly this manner.

REFERENCES

[1] Armstrong, M and D Flynn, M Hammond, S Jolly R Salmon, *High Frame Rate Television*, BBC Research Whitepaper WHP 169, September 2008.

[2] De Simone, F et al., *Towards high efficiency video coding: Subjective evaluation of potential coding Technologies*, Journal of Visual Communications (2011), doi:10.1016/j.jvcir.2011.01.008

[3] Ho, Yo-Sung and Jung-Ah Choi, *Advanced Video Coding Techniques for Smart Phones*, 2012 International Conference on Embedded Systems and Intelligent Technology (ICESIT 2012), Jan. 27–29, 2012.

[4] Howald, Dr. Robert L, Dr. Sean McCarthy, *Bits, Big Screens and Biology*, The Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[5] Howald, Dr. Robert L, *Boundaries of Consumption for the Infinite Content World*, SCTE Cable-Tec Expo, New Orleans, LA, October 20-22, 2010.

[6] Howald, Dr. Robert L, Dr. Sebnem Zorlu-Ozer, Dr. Nagesh Nandiraju, *Delivering Pixel Perfect*, The Cable Show Spring Technical Forum, May 11-13, Los Angeles, CA.

[7] Howald, Dr. Robert L, *Fueling the Coaxial Last Mile*, SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[8] Howald, Dr. Robert L, *Looking to the Future: Service Growth, HFC Capacity, and Network Migration*, 2011 Cable-Tec Expo Capacity Management Seminar, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Atlanta, Ga, November 14, 2011.

[9] Howald, Dr. Robert L, and Phil Miguelez, *Upstream 3.0: Cable's Response to Web 2.0*, The Cable Show Spring Technical Forum, June 14-16, 2011, Chicago, Il.

[10] Marpe Detlev , et al., *Video Compression Using Nested Quadtree Structures, Leaf Merging, and Improved Techniques for Motion Representation and Entropy Coding*, IEEE Trans. Circuits Syst. Video Technology., Vol. 20, Nr. 12 (2010) , p. 1676-1687.

[11] McCarthy, Dr. Sean T., *Quantitative Evaluation of Human Visual Perception for Multiple Screens and Multiple CODECs*, 2012 SMPTE Annual Technical Conference & Exhibition

[12] McCarthy, Dr. Sean T, *A Biological Framework for Perceptual Video Processing and Compression*, SMPTE Motion Imaging Journal, Nov/Dec 2010.

[13] McCarthy, Dr. Sean T., and W.G. Owen, "Apparatus and Methods for Image and Signal Processing," US Pat. 6014468 (2000). US Pat. 6360021 (2002), US Pat. 7046852 (2006), 1998.

[14] Sullivan, Gary J. and Jens-Rainer Ohm, *Recent Developments in Standardization of High Efficiency Video Coding (HEVC),* SPIE Applications of Digital Image Processing XXXIII, Andrew G. Tescher (editor), Proceedings of SPIE Volume 7798, Paper number 7798-30, August, 2010.

[15] Ulm, John and Gerry White, *Architecture & Migration Strategies for Multi-screen IP Video Delivery*, 2012 SCTE Canadian Summit, March 27-28, Toronto, CA.

[16] Wiegand, T, G. Sullivan, G. Bjontegaard, and A. Luthra, *Overview of the H.264/AVC Video Coding Standard*, IEEE Trans. Circuits Syst. Video Technol., vol. 13, no. 7, pp. 560-576, July 2003.

[17] Yamashita, Takayuki et al. "Super Hi-Vision" Video Parameters for Next-Generation Television, SMPTE Mot. Imag J. May-June 2012 vol. 121 no. 463-68 (NHK Science & Tech Res. Labs)

[18] ITU-T and ISO/IEC, ITU-T Rec. H.264 | ISO/IEC 14496-10 Advanced Video Coding (AVC), May 2003 (with subsequent editions and extensions).

[19] ISO/IEC JCT1/SC29/WG11 (MPEG), "Description of High Efficiency Video Coding (HEVC)," doc. no. N11822, Daegu, KR, January 2011.

[20] ISO/IEC JCT1/SC29/WG11 (MPEG), "Vision, Applications and Requirements for High Efficiency Video Coding (HEVC)", doc. no. N11872, Daegu, KR, January 2011.

[21] ITU-R BT.2020 (2012) Recommendation -"Parameter values for ultra-high definition television systems for production and international programme exchange"

[22] advanced-television.com

[23] CarltonBale.com

[24] http://mediacenter.motorola.com/Content/Detail.aspx?ReleaseID=15389&NewsAreaID=2&ClientID=1

[25] http://www.cable360.net/ct/webcasts/2010_08_10/

[26] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf

ADVANCED MENU USAGE AND SYSTEM ARCHITECTURE: IMPACTS ON USER BEHAVIOR

Jim Brown
Buckeye CableSystem

Carol Ansley, Scott Shupe
ARRIS Group

Abstract

For a long time CATV customers have had the infamous grid guide with its slowly rolling sequence to contend with if they wanted to know what was playing. Gradually as more intelligent settop boxes became available, the grid guide became scrollable. Recently, more advanced electronic program guides and menus have begun spreading through the operators' networks as well as newer device architectures that allow several devices to work together. This paper studies the behavior of users working with newer system architectures and more advanced user guides, and compares it to users of older standalone devices with traditional grid guides. This behavior can inform future network system design that has to consider the expectation for channel change time and relevant use cases.

CHALLENGES OF HOME NETWORK COMPLEXITY

Like most operators today, Buckeye has been looking for an enhanced user interface to improve subscriber interactions and support advanced guide features. The traditional user interface employing a grid programming guide has been around the industry for a long time now. This interface also includes many traditional guide features including folder access to Video on Demand and DVR content. The ability to customize this interface from the perspective of the operator or end user is sparse, usually limited to color scheme changes. Additionally, this platform has had relatively few interface enhancements over the past few years. However, there are enhancements planned for that user interface.

Subscribers are becoming accustomed to a broader range of enhanced user interfaces for access to programming and other content; this is the case to some extent with our current subscribers, but more particularly our next generation subscribers. The next generation subscriber will be looking for something other than a traditional grid guide and likely, something they can customize to some degree.

Our initial transition to a new user interface has been with the ARRIS Whole Home Solutions (WHS) system which offers a new system architecture and a different guide paradigm. The system architecture has a central gateway unit that provides content to several attached IP STBs. This new guide interface allows subscribers to choose their programming guide preference from a traditional grid guide or an enhanced rotational X and Y axis guide that provides an improved look and feel. The ARRIS WHS system also offers the opportunity to gain visibility in aggregate as to how subscribers are viewing and accessing content, providing valuable network, programming and marketing information that was ideal for this study.

STUDY BACKGROUND AND SETUP

With Buckeye's deployment of the ARRIS WHS system into existing hubs and nodes that also have substantial deployment of traditional single tuner settop boxes (STBs) and dual tuner DVR STBs, an environment exists that allows direct comparison of user activity on the different platforms.

The older STBs have a fairly standard grid-based guide interface as seen in Figure 1.

Figure 1 Grid Guide Menu Example

The standalone STBs are a mix of single tuner and dual tuner STBs that also provide Digital Video Recording services (DVR) with accompanying menus, such as **Error! Reference source not found.**.

Figure 2 Grid Guide DVR Menu Example

Figure 3 WHS Moxi Main Menu

The WHS system allows subscribers to choose their programming guide preference from a traditional grid guide or an enhanced rotational X and Y axis guide that provides a more web-like look and feel. The user is

Figure 4 WHS Live TV UI Example

presented with a series of icons or cards that can be used to select different activities. Within a card, the user sees a range of choices appropriate to that activity presented vertically. Examples can be seen in Figures 4 and 5.

STUDY METHODOLOGY

To monitor the channel change behavior of the subscribers, this study took advantage of several logging facilities within the Buckeye network. The anonymized logging facility of the switched digital video (SDV) system played an important part. The SDV system records all channel changes for all STBs in the system. This data allowed the users of all three types of devices to be compared directly through their channel change choices.

Within the ARRIS WHS system, extensive logs also allowed the study of user behavior with anonymized data. The scheduling and viewing of DVR recordings were also available for analysis, though this could not be compared to the wider population of older STBs.

Video on Demand use was also examined through the system's centralized usage logs. Since all VOD views for all subscribers are tracked, this facility allowed user behavior to be compared between the different types of STBs.

As mentioned earlier, an additional difference between the older STBs and the WHS system is that the WHS system uses a

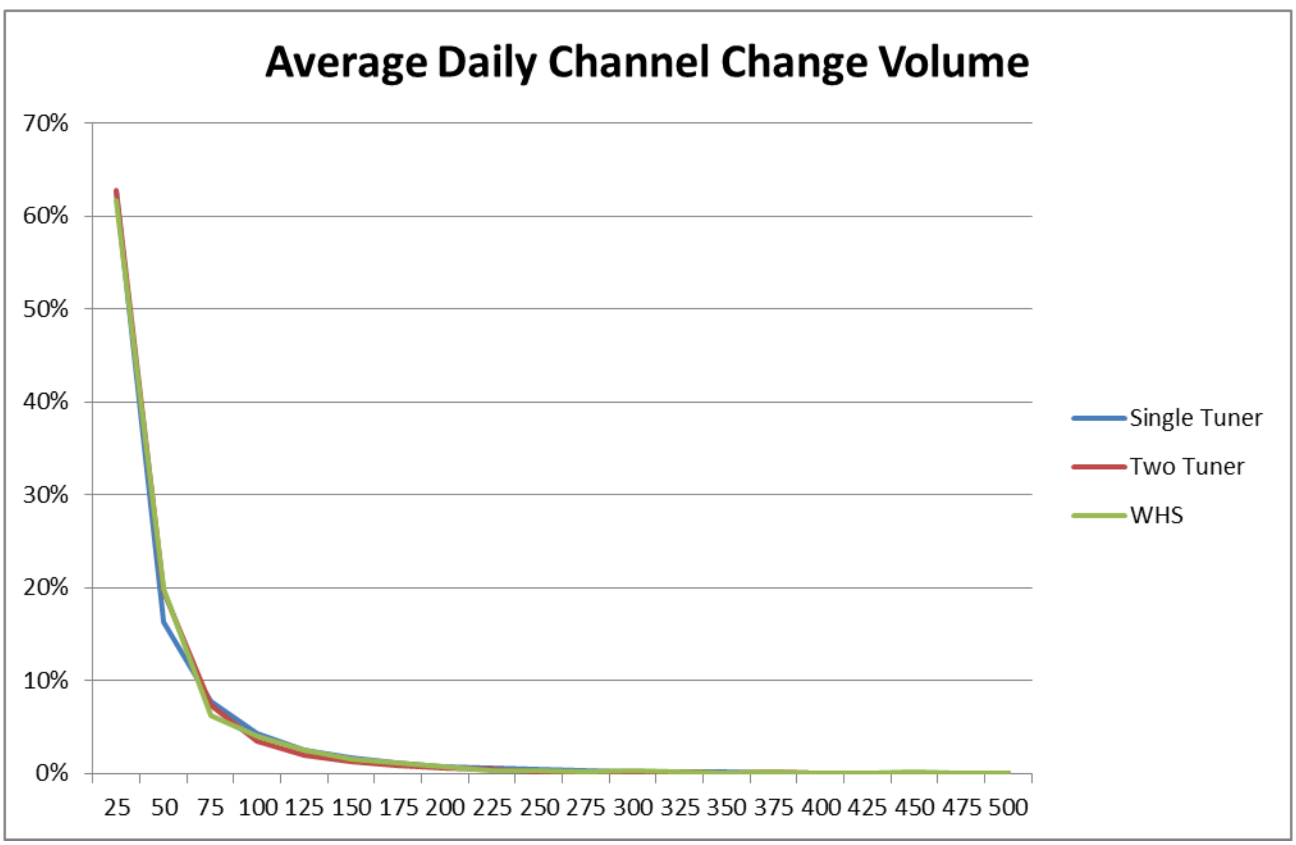

Figure 5 Daily Channel Change Volume Averaged Across One Week

Gateway with 6 shared tuners to support up to 6 IP STBs. Any DVR recordings are also shared across the household. Because the Gateway provides a WHS household's only interface into the SDV and VOD systems, the logs from those systems reflect the overall usage for that entire household. To allow a more direct comparison between the stand-alone STB usage and the WHS usage, internal logs from the WHS were used to gain visibility into how the users were interacting with the individual IP STBs.

RESULTS

Overall, the results showed the greatest differences in user behavior between DVR households and non-DVR households, and smaller differences that may be due to varying user responses to the different menu types or possibly to the WHS whole-home service architecture.

Channel Change Discussion

Looking at channel change findings, subscriber households using WHS had about the same number of channel change operations per household as a single tuner or dual tuner STB, the average number of operations per day was 37 for a WHS system and a single tuner STB, and 35 for a dual tuner STB. Figure 5 compares histograms of the average daily channel change performance of the three groups of STBs. The three groups of STBs have very similar distributions of channel change habits.

Finding that WHS households taken as a single unit are very close to single STB behavior was surprising, since a single WHS household has on average 2.6 IP STBs.

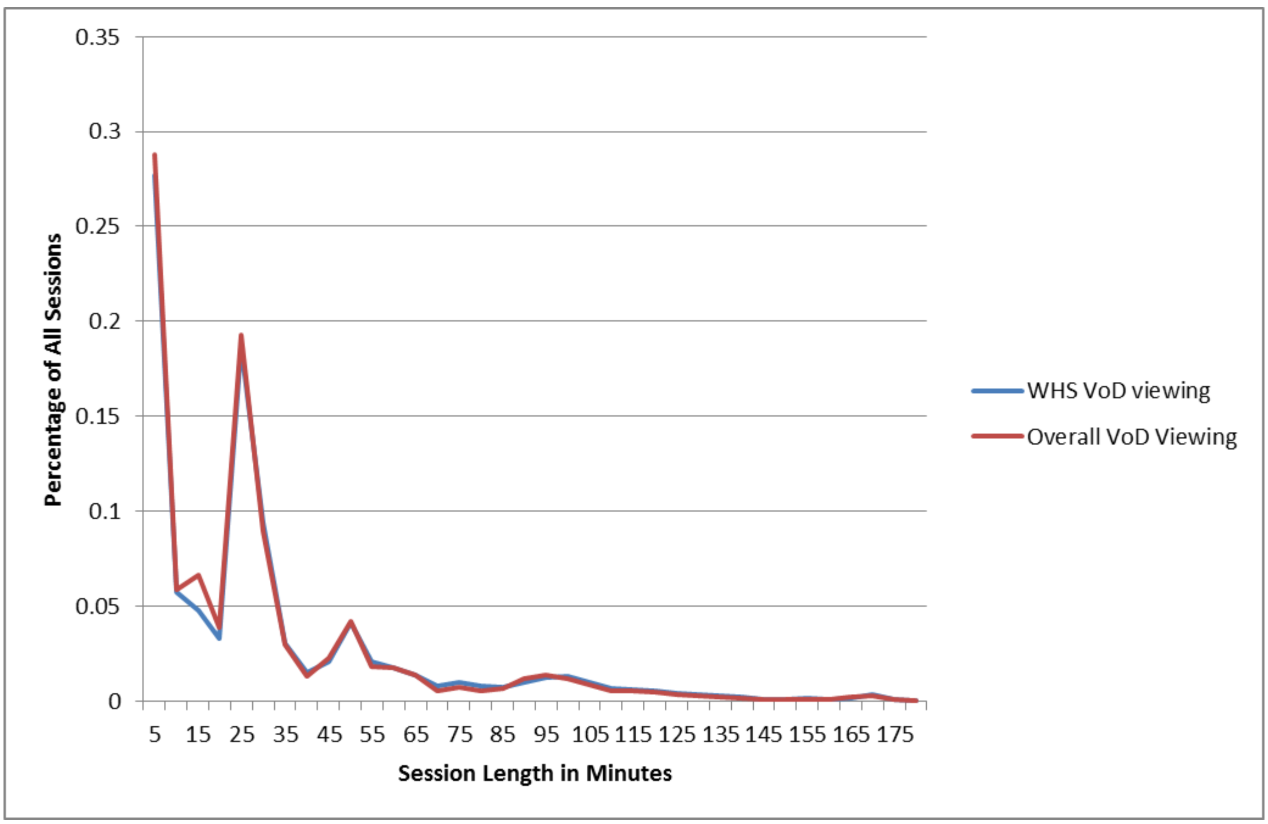

Figure 6 VOD Session Length Comparison

Checking for simultaneous operation of IP STBs during the week did show that even though the average number of STBs deployed per system was 2.6, only 1.5 IP STBs, on average, were active at any one time. This figure does not include channel changes that are driven by recordings. Overall, this suggests that users require fewer channel changes to find the programming they want. Since the dual tuner STBs can provide DVR functionality as can the WHS, it also suggests that at least part of the difference is due to the advanced menu structure of the WHS user interface as opposed to the presence of DVR functionality alone.

A related finding was that many subscribers using the WHS still used the grid-based guide occasionally, though as the time the user had spent with the WHS system increased, the use of the grid guide version decreased.

VOD Discussion

Video On Demand (VOD) statistics were also studied to see if the different menus and system architectures were correlated with any VOD behavior changes. Overall, as seen in Figure 6, the users of the advanced menu system behaved similarly to other users. WHS users considered as a household were slightly more likely to view at least one VoD asset, 32%, versus 31% of dual tuner units and 26% of single tuner STBs. But since there are multiple STBs in the subscriber's residence, this suggests that the total VoD

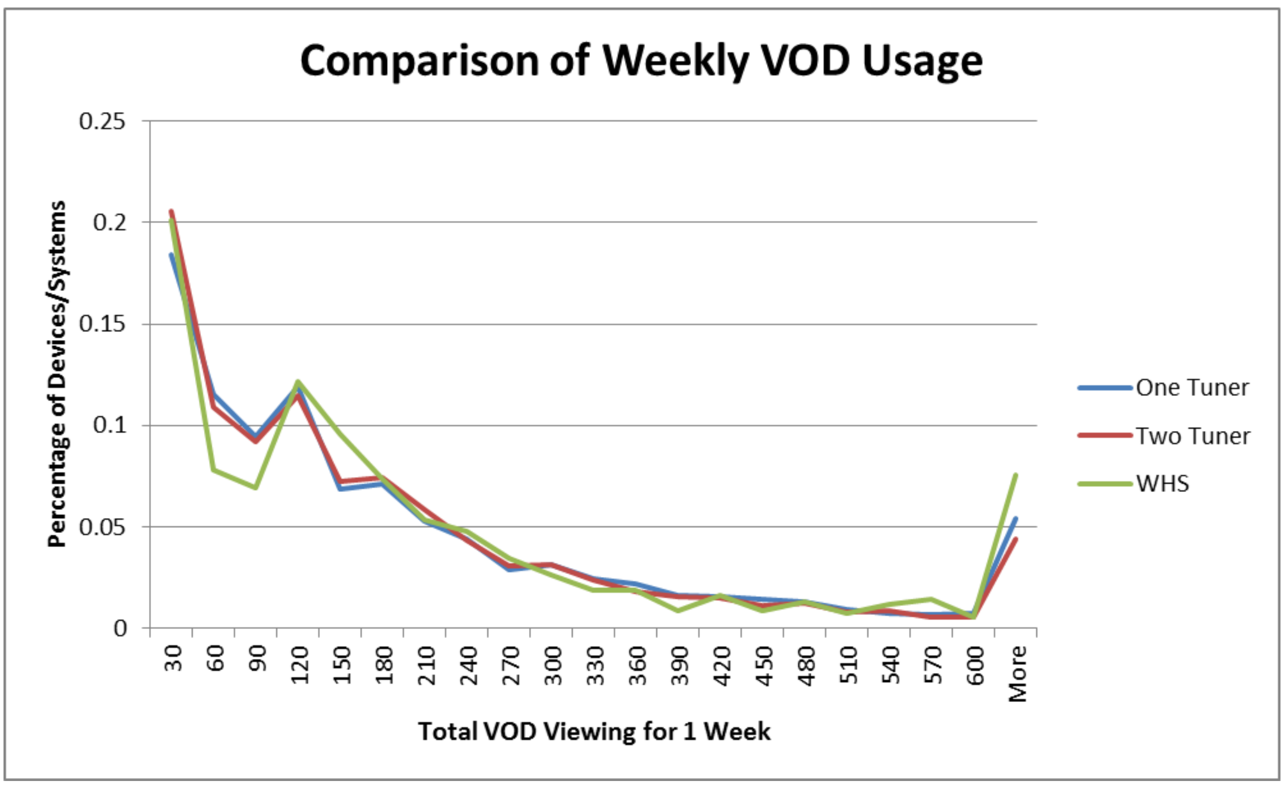

Figure 7 Comparison of VOD Usage

purchases per household may be close to the same or slightly less than households with the older technology. Again even though the household average was 2.6 IP STBs, an average of 1.45 STBs per household were actually used for orders. The total consumption, shown in Figure 7 was also quite similar.

One interpretation of the weekly usage metrics is that about 10% of all subscribers order a 2 hour movie on VOD once a week. About 30% of all subscribers view substantially more content than that, with WHS subscribers leading in the high usage category at 36%.

DVR Discussion

Comparing the single tuner boxes with the DVR capable systems does show that DVR capabilities tend to alter viewing habits. Since the WHS system has extensive DVR logs it was possible to examine how people with that system tend to use it.

DVR usage alters the channel change statistics in two ways. Recordings are scheduled around when programming is available – so some channel changes reflect the DVR engine seizing a tuner and selecting the channel directly. The viewing times are set by when the subscriber is available to view the content, and the

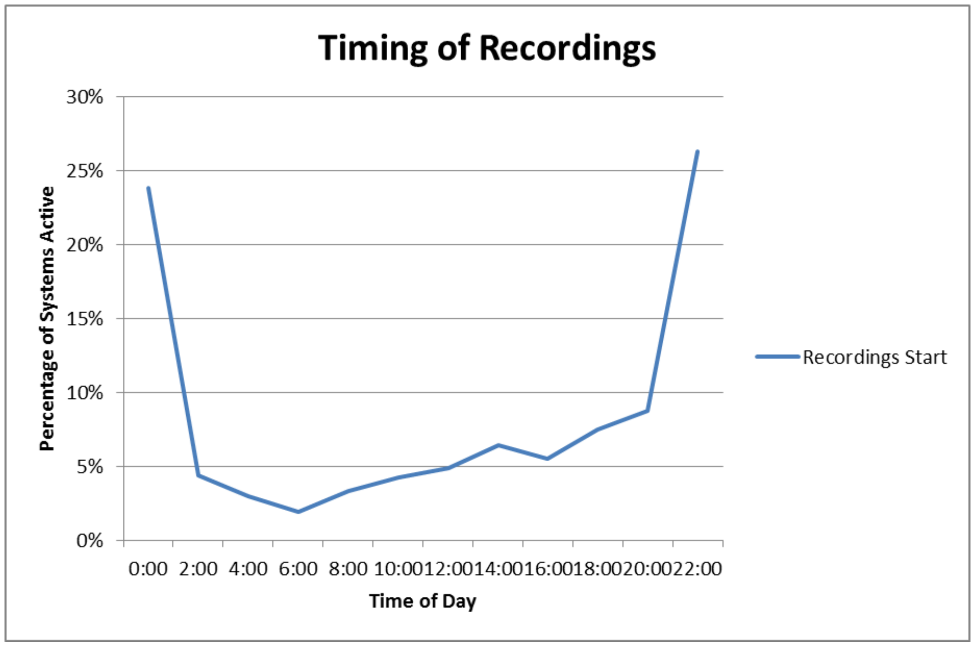

Figure 8 Timing of Recording Sessions

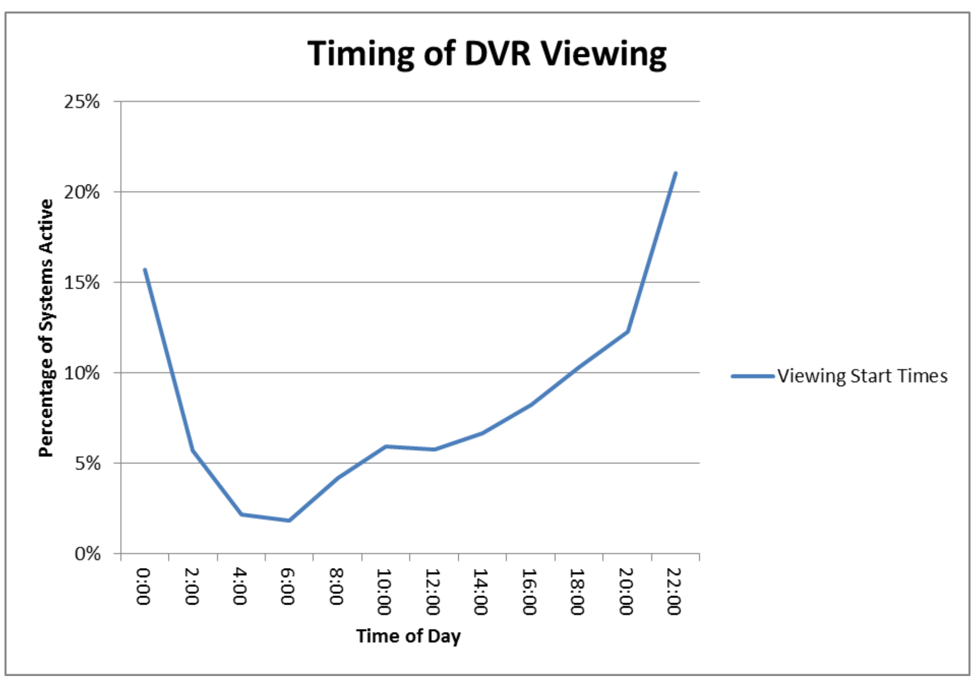

Figure 9 Timing of Viewing Session

subscriber may or may not also view live programming before or after they view the stored content.

By comparing when content was recorded (Figure 8) with when the recorded content was usually consumed (Figure 9), it can be seen that the timings of those activities is very similar. The bulk of recordings does occur slightly later in the evening than the bulk of the viewing sessions. Comparing the number of recording with the number of viewing sessions also highlights the fact that only about 70% of recording sessions are actually viewed. Overall, when considering that DVR systems tend to have fewer channel change operations than non-DVR units, it suggests that DVR users are using the DVR capabilities to more efficiently control their viewing time.

IN SUMMARY

Operators and guide developers are placing significant focus in the area of the user interface. The push to enhance this experience for their subscribers as well as to augment product usage and broaden their own marketing data must continue to ensure that the operators' video products are accessible and attractive to the next generation of subscribers.

Looking at detailed usage logs offers insights into how subscribers are actually using a product, and can turn up unexpected insights, such as that a whole home DVR system tends to greatly reduce the actual network channel change traffic compared to the deployment of a similar number of stand-alone devices. This study also showed that subscribers will continue using familiar technologies, such as a grid guide, but will also gradually accept new user interfaces.

Further research into user behavior will undoubtedly turn up fresh insights allowing the operators to tune their networks to actual user behavior and provide a great user experience at the same time.

Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS

Tom Cloonan, Jim Allen, Tony Cotter, Ben Widrevitz
ARRIS

Abstract

This paper will discuss reasons why current methods of DOCSIS Quality of Experience (QoE) Monitoring are becoming obsolete. It will be shown that the monitoring of channel utilization alone can often lead to erroneous conclusions.

The paper will discuss the impact of Service Level Agreement (SLA) settings on QoE monitoring, offering two different philosophies that an MSO can choose to follow (SLA-agnostic vs. SLA-dependent). It will also be shown that QoE monitoring tools and CMTS scheduling algorithms may benefit when more tightly coupled with adaptive CMTS scheduling algorithms that can utilize monitoring tool outputs and with monitoring tools becoming more cognizant of the CMTS scheduling philosophies for different SLAs.

The paper will identify the various traffic metrics that can be utilized for QoE monitoring, and it will then discuss the fact that different traffic types are sensitive to different traffic metrics and have different thresholds determining acceptable QoE Levels.

Finally, the paper will describe and illustrate necessary attributes of second-generation DOCSIS QoE Monitoring tools.

BACKGROUND INFORMATION

Introduction and Motivation

Quality of Experience (QoE) is a basic measure of the user's level of satisfaction, and Quality of Experience Monitoring is an extremely important tool that helps MSOs reduce subscriber churn and that helps MSOs quickly trouble-shoot subscriber problems. QoE Monitoring permits the MSO to determine when their network infrastructure is providing adequate bandwidth support for the offered services. Low QoE scores can be an indication that the network is no longer performing at adequate levels, and the existence of low QoE scores usually serves as an important trigger to the on-going evolution and modification of the HFC Plant as MSOs continually schedule node-splits and/or channel augmentations to address observed QoE issues.

Unfortunately, QoE Monitoring is becoming more difficult, because the nature of traffic propagating over DOCSIS networks has been changing quite rapidly in recent years. These changes are taking place on four different fronts.

First, there is a much larger mix of traffic types that are now making up the aggregated bandwidth within typical DOCSIS networks. These traffic types now include:

Web-Surfing
Over-the-Top SD Streaming IP Video
Over-the-Top HD Streaming IP Video

MSO-Managed SD Streaming IP Video
MSO-Managed HD Streaming IP Video
Streaming Audio
Gaming
VoIP
Peer-to-Peer Services

The complex interactions between these different traffic types is making it more and more difficult to predict when different subscribers using different applications are receiving adequate service levels.

Second, each of the different traffic types listed above has its own unique and different traffic requirements and sensitivities. Some of the traffic types are sensitive to changes in packet stream bandwidth. Others are sensitive to changes in packet delay and jitter. Still others are sensitive to packet loss. The unique requirements for each traffic type make it difficult to easily ascertain QoE levels.

Third, Maximum Sustained Traffic Rates for cable modem subscribers now range over much larger values. The heterogeneous mix of subscribers with different bandwidth needs also complicates the problem of determining QoE levels.

Fourth, the arrival of Adaptive Bit-Rate IP Video traffic on the DOCSIS network introduces a new traffic type that can quickly expand to fill unused bandwidth capacity and can be compressed during periods of congestion. This also makes it more difficult to determine QoE levels.

While being able to quickly and correctly determine the DOCSIS QoE levels for individual subscribers is getting more difficult, it is also becoming more imperative. Why? There are several reasons. First, more and more service types carried by DOCSIS might be considered to be “mission critical” services. VoIP is an obvious example, but even IP Video services delivered to television sets might be considered critital in the future. In addition, MSOs are entering a decade of rapid change for many of their service bandwidths, which will cause rapid changes in their HFC spectrum mixes. Each MSO will experience these changes differently, as each MSO will be working with a different set of conditions and constraints on their HFC plants that lead them to different Spectral Maps over time. **Fig. 1** illustrates one example of how one MSO might migrate their spectral maps over the next decade. [Clo1]

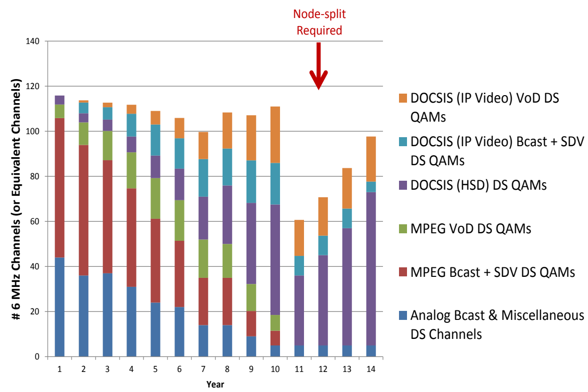

Fig. 1- Example Spectral Map Changes

This figure illustrates how quickly certain service types are expected to grow. Required per-subscriber bandwidth growth for DOCSIS High-Speed Data service may continue to grow at a 50% CAGR in the future, and that would force the DOCSIS service tier to consume much larger portions of the HFC spectrum. The introduction of MSO-Managed IP Video over DOCSIS within the next decade is also expected to force swift and considerably larger changes on the HFC Spectral Map.

As MSOs increase the number of DOCSIS channels, they will from time to time be forced to make room for the new DOCSIS channels by eliminating channels from other service types within the HFC Spectral Map. This fact leads to a difficult dilemma- which services should donate their channels to DOCSIS and how quickly can they make these changes. There will undoubtedly be times when MSOs are forced to delay growth in their DOCSIS service tier due to constraints that temporarily preclude them from removing channels from the other tiers. Thus, MSOs may be forced to temporarily “squeeze” their DOCSIS services into over-subscribed DOCSIS spectra from time to time. This undesirable condition will make it imperative for MSOs to be able to determine the QoE of their DOCSIS subscribers while going through the upcoming decade of challenging transitions.

In addition, the “resolution” of QoE Monitoring performed by MSOs must be increased. Higher resolution QoE Monitoring implies the ability to monitor the QoE levels for different traffic types as well as for different Service Level Agreements (SLAs) at higher sampling rates. If DOCSIS subscribers are going to experience QoE issues, MSOs will undoubtedly be interested in knowing exactly which subscriber traffic types and SLAs are experiencing the issues and they will also be interested in knowing how severe the issues are for each of the traffic types and each of the SLAs. This could permit MSOs to make intelligent decisions about whether the issues warrant drastic changes like instantaneous node-splits or whether subscribers are likely to “live with” the issues for a short period of time.

As an example, if the only traffic type and SLA experiencing a performance issue is Peer-to-Peer traffic within the Bronze SLA level and the MSO has only a small percentage of subscribers within that Bronze SLA level, then it is possible that the MSO may want to risk living with the issue for a short while. If, on the other hand, the traffic types that are experiencing the issue are Web-Browsing and VoIP and IP Video services within the Bronze and Silver and Gold SLA levels, then rapid actions to mitigate the issues may be required.

As a result, it should be clear that high-resolution QoE Monitoring of the performance levels of different traffic types and different SLA levels could become an essential tool in the MSO toolkit as MSOs navigate the challenging, ever-changing waters of the upcoming decade.

<u>Previous QoE Monitoring Tools</u>

In the past, QoE Monitoring has taken on many different forms for the different services that were being monitored.

For VoIP services, monitoring was often implemented within the VoIP destination endpoints, monitoring performance parameters such as packet delay, packet jitter, and packet loss. In addition, VoIP monitoring oftentimes kept track of connection set-ups and the success of the associated signaling.

For Legacy Digital Video services, monitoring was often implemented with sampling points along the path of the Digital Video streams. These sampling points could be positioned in the headend and/or in the subscriber endpoints (such as STBs). These sampling points could monitor packet loss, packet jitter, packet delay, and packet corruption. They could also perform deep-packet inspection types of functions to determine what type of video packet was corrupted or lost.

For both VoIP and Video monitoring, the monitoring tools would collect all of the aforementioned parameters. The parameters would then be combined together to create a reasonable prediction of the user’s QoE level. Many of the monitoring tools had complex algorithms that determine the magnitude of any service disruption. For example, Video monitoring tools could determine whether a packet loss would corrupt an I-Frame or a P-Frame or a B-Frame to determine the duration of the resulting error. They also would attempt to determine if packet corruption would lead to macroblock tiling displays or frozen video displays or audio outages. The magnitude and duration of service interruptions were taken into account when creating these QoE level estimates.

The QoE level estimates from monitoring tools would oftentimes list all of the acquired parameters from above. However, a convenient technique for displaying the overall QoE level estimate with a single number uses the concept of a Mean Opinion Score (or MOS Score). A MOS Score is an average score (ranging from 1 to 5) that attempts to combine the many effects from the many different performance-affecting parameters defined above. A MOS Score of 1 corresponds to the lowest possible quality level as it might be perceived by the subscriber. A MOS Score of 5 corresponds to the highest possible quality levels as it might be perceived by the subscriber. **Table 1** illustrates the typical definitions for the different MOS Score levels.

MOS	Quality	Impairment
5	Excellent	Imperceptible
4	Good	Perceptible but not annoying
3	Fair	Slightly annoying
2	Poor	Annoying
1	Bad	Very annoying

Table 1- MOS Score Level Definitions

Since QoE levels (and their associated MOS scores) can vary over time, the continual monitoring of QoE MOS scores is important for MSOs, because the existence of lower MOS scores can be an indication that the network is no longer performing at adequate levels, and the existence of low MOS scores usually serves as an important trigger to the on-going evolution and modification of the HFC Plant as MSOs continually schedule equipment upgrades and/or node-splits and/or channel augmentations to address observed QoE issues.

MSOs have also performed QoE monitoring of their DOCSIS services for many years. These DOCSIS QoE Monitoring efforts were not as well-developed as the techniques used for VoIP and Digital Video services. For the most part, DOCSIS QoE Monitoring of packet streams was limited to a study of bandwidth utilizations on DOCSIS channels or DOCSIS service groups. Some MSOs would also monitor channel parameters such as SNR, uncorrectable FEC errors, and correctable FEC errors to determine if the quality of the DOCSIS channels was degrading with time. These simple techniques were often employed because they were good enough to ascertain a rough QoE metric for most of the DOCSIS services. They worked well in the previous era when less user applications existed, when users were more homogeneous, and when the distance between the highest and lowest Maximum Sustained Traffic Rate levels were much smaller than they are today. Some QoE monitoring tools would also provide more advanced Degraded Modem-Hour metrics (which are similar to MOS scores) for DOCSIS services. These types of scores were very good at predicting subscriber QoE levels when the number of different traffic types propagating over the DOCSIS network were limited.

However, with the rapid expansion of new traffic types being carried by DOCSIS and with larger variances in Maximum Sustained Traffic Rates for subscribers, it is becoming apparent that extensions may be required to the previously-available DOCSIS QoE Monitoring tools. Why is this the case?

Let us consider an example. Assume that an MSO has deployed a single ~40 Mbps DOCSIS channel with 100 DOCSIS subscribers sharing the channel, and assume that the channel is operating at ~100% utilization. Assume that there is 0.001% packet loss. Are the subscribers currently satisfied with their Quality of Experience levels?

The answer is not clear. For example, if there is currently only one active subscriber using the channel and that subscriber has a Maximum Sustained Traffic Rate of 40 Mbps, then that particular subscriber would be receiving the full ~40 Mbps of service and would probably be ecstatically satisfied with the level of the service. Since the other 99 subscribers are apparently not using the DOCSIS channel at this point in time, they are probably satisfied as well.

However, in the unlikely event that all 100 DOCSIS users were simultaneously sharing the resources of the channel at the same time, then (on average) each subscriber would be receiving (~40 Mbps)/100 = 400 kbps of bandwidth. Depending on the applications that the subscribers are using, the subscribers may or may not be satisfied. If everyone is making VoIP telephony calls requiring 120 kbps of bandwidth, then the 400 kbps level of bandwidth may be adequate and everyone might still be satisfied. If, however, a lot of the active subscribers are trying to access HD IP Video streaming services requiring (say) 6 Mbps of bandwidth, then 400 kbps level of offered bandwidth is inadequate and will undoubtedly result in many unsatisfied subscribers.

One of the other inadequacies of existing DOCSIS QoE Monitoring methods is the frequency with which monitoring is oftentimes implemented. Many DOCSIS QoE Monitoring systems of today are designed with typical sampling periods on the order of days or hours or minutes. While that gives some level of visibility into the QoE levels of subscribers, it can miss some of the high-frequency, short-duration events that might cause QoE degradation on the DOCSIS network (like micro-bursts of bandwidth causing frequent and temporary increases in latency).

Thus, from the above scenarios, it becomes clear that, in the future, MSOs may need to do more than just measure the channel utilization of their DOCSIS channels once per hour to ascertain whether their subscribers are currently satisfied or not. The remaining sections of this paper will outline some new ideas that might be useful as we consider the addition of extensions into the future DOCSIS QoE Monitoring tools.

IMPROVEMENTS IN FUTURE DOCSIS QoE MONITORING TOOLS

Categorizing DOCSIS QoE Metrics

As indicated in the previous sections, DOCSIS QoE Monitoring tools are likely to be receiving some upgrades over the next decade. In order to determine the likely forms of these upgrades, it is beneficial to first consider the various metrics of the DOCSIS packet streams that are already monitored and determine how they can be classified. These classifications will prove to be useful as we define new metrics.

The existing metrics used in QoE Monitoring tools and the methods for collecting those metrics can be classified using several different attributes.

First, metrics can be collected in a continuous fashion (ex: sampling every time a packet passes) or in a periodically sampled fashion (ex: sampling average values once every hour).

Second, metrics can be direct measurements of the Quality of Experience of a subscriber (ex: measuring file download times) or can be indirect measurements of the Quality of Experience of a subscriber (ex: measuring packet delay and trying to infer subscriber satisfaction for different applications when those delay conditions exist).

Direct methods can place "canary clients" inside of already-deployed modems or can place purpose-built "canary modems" out on the HFC Plant. These canaries periodically download HTTP files like a Web Browser or download video content like an IP Video client or perform many parallel TCP downloads of small files like a Peer-to-Peer client or transmit small UDP packets like a VoIP client or access DNS servers and measure the performance metrics of these activities. The direct performance metrics for canaries include metrics such as file download times, packet loss (for VoIP), packet delay and jitter (for VoIP), and delay (for DNS). These metrics are directly translatable into MOS scores that might be seen by other real subscribers on the same channel that might be using that same application.

Indirect methods do not require canaries and do not add any extra traffic to a channel that might already be congested. Instead, indirect methods sit passively inside of network elements like CMTS and CMs and measure the performance of "real-world" packets that are propagating through those devices. Indirect performance metrics that might be measured include packet throughput (for a user), packet loss, packet delay, packet jitter, and channel congestion. These metrics do not directly translate into MOS scores, but with intelligence, a QoE Monitoring tool can oftentimes infer the QoE levels of subscribers that might be on a channel where these indirect performance metrics were captured.

In general, direct metrics tend to be better measures of a subscriber's QoE level, because they measure "real-world" response times for their simulated traffic loads. However, direct metrics are more difficult to obtain (requiring canaries to be planted in the HFC plant and requiring servers to serve up the requested files). In addition, direct methods add more bandwidth to channels that may already be congested, so these measurement techniques can negatively affect the channel performance levels.

In general, the inference step required for the use of indirect metrics tend to make them be less accurate measures of a subscriber's QoE level, because they can only infer what real performance levels might be from the measured variables. However, indirect metrics are much easier to obtain (requiring simple counts to be implemented in CMTS or CMs to monitor the real-world data that passes by). In addition, indirect methods do NOT add any extra bandwidth to channels that may already be congested, so these measurement techniques do not negatively affect the channel performance levels.

As a result of the above facts, QoE metrics can be classified into one of four categories:

1) Continuous, Direct QoE metrics
2) Sampled, Direct QoE metrics
3) Continuous, Indirect QoE metrics
4) Sampled, Indirect QoE metrics

In addition to this categorization of QoE metrics, the authors have also found it valuable to identify the scope of a particular metric. The scope of a metric helps answer the question "What layer or level of the data stream was monitored to collect the metric?" To define the scope levels, it is beneficial to view the aggregated traffic that makes up a DOCSIS network to be comprised of hierarchical layers of data streams, where each layer is built up from a set of elements from the underlying layers. An abstract illustration of this concept is shown in **Fig. 2**, where DOCSIS Service Groups are comprised of DOCSIS Channels, DOCSIS Channels are comprised of Modem Traffic, Modem Traffic is comprised of Service Flow Traffic, Service Flow Traffic is comprised of TCP/UDP Sessions, and TCP/UDP Sessions are comprised of IP Packets.

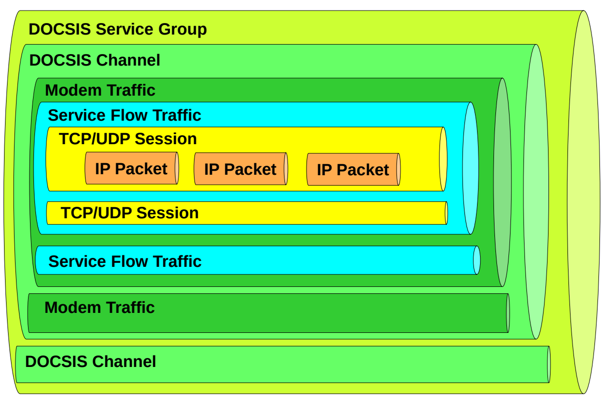

Fig. 2- Traffic Hierarchy

Monitoring of QoE can take place at various levels in the hierarchy of **Fig. 2**. For example, at the lowest level of the hierarchy (the IP Packet level), one can monitor packet-level characteristics such as packet loss, packet latency, and packet jitter. At the next level of the hierarchy (the TCP/UDP Session level)- which often corresponds to an HTTP Session for Web-Browsing or IP Video applications- one can monitor session-level characteristics such as average throughput, average download times, average packet loss, average packet latency, and average packet jitter within the session. All of these average session level characteristics can also be calculated for the higher levels in the hierarchy. In addition, at the Service Flow level, Modem level, DOCSIS Channel level, and DOCSIS Service Group level, one can also calculate aggregate characteristics such as Capacity Utilization.

Each of these metrics can of course be categorized using the categories defined above. For example, average packet jitter that is sampled once per hour within a particular Service Flow would be categorized as a Sampled, Indirect QoE metric at the Service Flow scope. It is obviously sampled at a rate of one sample per hour. It is an indirect metric, because intelligence and calculations are required to determine whether the user experience level will be good or bad as a result of the measured jitter level, and it should be clear that the traffic type will be a key factor in determining that fact. It is also sampled from the many packets with the scope of a Service Flow.

All of the above metrics (packet loss, packet latency, packet jitter, throughput, capacity utilization, etc.) viewed at the different scope levels of the hierarchy are typical metrics that can be useful in developing a view of the subscriber QoE level at any instant in time. These metrics have been used by many traditional QoE Monitoring tools.

But the authors were interested in determining if there were other metrics that could be added to the list of existing metrics to help create a more insightful view of the subscriber QoE level in the future. The remainder of this paper will focus on this important topic.

<u>Areas For Possible Improvement Within DOCSIS QoE Monitoring</u>

The existing metrics available for QoE Monitoring tools cover many aspects of DOCSIS packet stream performance. But there is still much room for improvement in the future. Several areas are beckoning for improvement. These areas will receive special attention within this paper. They include:

1) QoE MOS scores that are cognizant of the performance requirements of different traffic types
2) QoE MOS scores that are cognizant of the performance expectations of different SLAs
3) QoE MOS scores that are cognizant of both traffic types and SLAs
4) QoE MOS scores that are cognizant of the number of active subscribers sharing a channel or service group
5) QoE MOS scores that are cognizant of the traffic scheduling algorithms employed by the CMTS Downstream and Upstream traffic managers
6) QoE MOS scores that are cognizant of Offered Loads
7) QoE MOS scores that utilize metrics collected with high sampling rates
8) QoE MOS scores that are cognizant of different scopes within the DOCSIS Network

We will explore each of these areas in the sections below.

<u>MOS Scores Based On Traffic Types</u>

At a minimum, MOS scores of the future will undoubtedly be required to recognize that different traffic types have different sensitivies to different traffic attributes. In addition, different traffic types can have extremely different performance requirements and can also have very different thresholds marking the important boundary line between acceptable QoE performance and unacceptable QoE performance.

As an example, if we consider the primary traffic types outlined in the first section of this paper, we will find that each traffic type is sensitive to a different set of traffic attributes. This is outlined in **Table 2**.

Traffic Type	**Primary Sensitivities**
Web-Surfing	Avg BW, Avg Delay
OTT SD IP Video	Avg BW
OTT HD IP Video	Avg BW
MSO SD IP Video	Avg BW
MSO HD IP Video	Avg BW
Streaming Audio	Avg Jitter
Gaming	Avg Delay, Avg Jitter
VoIP	Loss, Avg Delay, Avg Jitter

Table 2- Traffic Attributes To Which Each Traffic Type Is Sensitive

One may argue that each of the traffic types in **Table 2** might also be sensitive to other attributes, and that would be an accurate argument.

For example, one may argue that VoIP is also sensitive to Average Bandwidth levels- it requires the Average Bandwidth to support the minimum levels required for VoIP traffic (ex: 120 kbps) and it will not operate well at lower levels of Average Bandwidth. But these levels are so low that they are usually guaranteed by most DOCSIS Networks (unless exceptionally high congestion is occurring), so for most VoIP applications in today's networks, the Average Bandwidth level is not a major issue for VoIP.

As another example, one may argue that IP Video is also sensitive to Average Delay and Average Jitter- that widely varying delays could lead to buffer underflows at the video display device. This is true if the IP Video Delivery Architecture does not buffer a large amount of video content before beginning the playout of the IP Video. However, most IP

Video Delivery Architectures of today do in fact buffer a large amount of video content before beginning the playout, so this problem with Delay and Jitter is actually quite rare.

Nevertheless, these two examples illustrate that it is not always clear how to identify the specific traffic attributes to which a particular application is sensitive. Any assumptions will be valid for some cases and potentially invalid for other cases. So one must be cautious when interpreting the meaning of MOS scores that are defined for different traffic applications.

Other errors in MOS score predictions can occur whenever multiple traffic types are flowing to a single modem at the same moment in time. That particular condition can create unexpected interactions between the different traffic types. For example, TCP congestion windows and TCP bit-rates for a newly-started flow may take a longer period of time to grow if multiple traffic types are sharing the bandwidth capacity to a single modem. These types of effects make it more difficult to correctly predict the MOS scores for the interacting traffic types.

Regardless of the aforementioned issues, most MSOs will agree that some indicators of performance level are better than no indicators of performance level- even if those performance indicators have limitations.

For this reason, we will attempt to define appropriate formulae for MOS scores for the currently-dominant traffic types found on DOCSIS Networks. Those formulae will be developed in one of the sections below.

<u>MOS Scores Based On SLAs</u>

Service Level Agreements (SLAs) typically exist between subscribers and service providers. These SLAs define the maximum and minimum performance levels that the subscriber can expect to see for their purchased subscription to the service. DOSCIS permits MSOs to specify many different performance attributes for each of the service flows attached to each of their subscribers. These performance attributes can include parameters such as the Maximum Sustained Traffic Rate (aka Tmax) and the Minimum Reserved Traffic Rate (aka Tmin). MSOs usually offer different levels of service to their subscribers using different service tiers with different price tags and different performance attributes. As an example, an MSO may offer a Gold Service Tier for $70, with a Tmax of 50 Mbps, and may also offer a Bronze Service Tier for $40, with a Tmax of 10 Mbps. Once these parameters are defined and configured for a particular subscriber, the MSO must try to satisfy the requirements of the SLA to ensure that the subscriber remains satisfied.

As a result, it seems apparent that QoE MOS scores could include information about the different SLA levels within the different Service Tiers. As an example, assume that an HD IP Video stream requires a minimum Average Bandwidth of 14 Mbps to yield acceptable results to the subscriber. If a subscriber has subscribed to the Gold Service Tier level (Tmax = 50 Mbps) and does not receive adequate bandwidth to support the required 14 Mbps HD IP Video stream, then that subscriber will obviously be less than satisfied with the service. The MOS score for the Gold subscriber will obviously be marked as being low. If, on the other hand, a subscriber has subscribed to the Bronze Service Tier level (Tmax = 10 Mbps) and does not receive adequate bandwidth to support the required 14 Mbps HD IP Video stream, then that subscriber has no right to be dissatisfied with the service, because their lower level of bandwidth is directly caused by their low service tier level. As a result, the MOS score for the Bronze subscriber may be marked as being low- but it should be marked as being low due to inadequate Tmax levels and not due to congestion. Alternatively, the MOS score can be marked high since the problem is not being caused by the MSO. Either approach can be deemed acceptable depending on how the MSO wants to interpret the information.

Thus, it should be apparent that MOS score calculations can make good use of any information provided on the subscribers' SLA levels. This means that every customer's QoE will now depend on their SLA and upon what service they are trying to obtain (e.g., Web-Browsing, IP Video viewing). Users with a higher tier SLA (and higher Tmax) will likely be happier than users with a lower tier SLA (and lower Tmax). Users requesting lower bandwidth services will tend to be happier than those needing more bandwidth. One might imagine every combination of SLA and traffic type receiving a different QoE. This type of scenario will be outlined below.

MOS Scores Based On Traffic Types And SLAs

The previous two sections described how MOS score calculations can benefit from information on the traffic types and SLAs being used by subscribers. In fact, MOS score calculations can make even better predictions about subscriber QoE levels if they combine these two attributes. One can imagine a good MOS score that displays all possible combinations of traffic types and SLAs.

Thus, if there are two SLA levels (Bronze and Gold) and if there are three primary traffic types (Web-Browsing, HD IP Video, SD IP Video), then one could envision a table with the three traffic types along the right-most column and with the two SLA levels along the top row. This results in six different combinations and six different MOS scores that can be calculated and displayed, as shown in **Table 3**.

	Gold Tier	Bronze Tier
Web-Browsing	MOS_1	MOS_4
HD IP Video	MOS_2	MOS_5
SD IP Video	MOS_3	MOS_6

Table 3- Traffic Type & SLA Combinations

MOS Scores Based On Active Subscriber Counts

In an earlier section of the paper, it was shown that knowledge of the channel utilization alone is inadequate to identify the QoE level for subscribers. In particular, it was shown that one cannot easily determine if a single ~40 Mbps DOCSIS channel with 100 DOCSIS subscribers sharing the channel and a channel utilization of ~100% yields high QoE levels or low QoE levels. If there is only one active subscriber, then the QoE level would likely be high (for that subscriber). If there are 100 active subscribers, then the QoE level would likely be low.

As a result, it should be clear that a measure of the number of active subscribers could also prove useful within MOS score calculations. For example, if the bandwidth is equally shared by all of the active subscribers, then the MOS score could use the following simple formula as an estimate of the bandwidth consumed by each subscriber:

$$\textbf{Avg BW/Sub} = \mathbf{C * U / N} \qquad \textbf{(1)}$$

where:
C = **Channel Capacity**
U = **Channel Utilization**

N = No. Active Subscribers

MOS Scores Based On CMTS Scheduling Algorithms and SLAs Of Active Subscribers

The Average BW per Sub formula (formula 1) above assumes that the bandwidth was shared equally by all of the active subscribers. This assumption may not always be true, because the details behind any CMTS scheduler might be quite a bit more complicated. As an example, assume that a particular CMTS scheduling algorithm is designed to give each active subscriber an amount of bandwidth that is proportional to their Minimum Reserved Traffic Rate (Tmin) value divided by the sum of all of the Tmin values for all of the active subscribers. The bandwidth offered by the CMTS scheduling algorithm to a particular subscriber's Service Flow (ABW) would therefore be given by:

$$\mathbf{ABW = C * U * Tmin(i) / \Sigma[Tmin(j)]} \qquad \mathbf{(2)}$$

where:

ABW = Avg Service Flow BW for Sub #i
C = Channel Capacity
U = Channel Utilization
Tmin(i) = Tmin for i-th Subscriber Flow
Σ[Tmin(j)] = Sum of Tmin(i) over Active Subs

If the QoE Monitoring tool was cognizant of the fact that the CMTS scheduling algorithm was using this formula, then the same formula could be used by the QoE Monitoring tool if/when it needs to determine the average bandwidth being offered to a particular subscriber within its MOS score calculations. In a certain sense, formula (2) is a more accurate description of the bandwidth offered to subscribers than the more simple version in formula (1). Using the more accurate formula (2), the SLA-based MOS score thresholds described in previous sections could be compared against the Average Bandwidths that would be offered by the CMTS scheduling algorithm, as predicted by formula (2).

It is important to note that these MOS score calculations would have to be altered if the CMTS scheduling algorithms were changed or if they were associated with different CMTSs from different vendors (since different CMTS vendors will likely use different scheduling algorithms within their proprietary solutions).

Thus, support for MOS scores of this nature requires tight coordination between the CMTS scheduling algorithms and the QoE Monitoring tool. In particular, the QoE Monitoring tool algorithms must be written with complete knowledge of the details behind the CMTS scheduling algorithms on each of the CMTSs that the tool would support.

MOS Scores Based On Requested Loads

While DOCSIS channel utilization is indeed a valuable tool that can (and should) be used in QoE Monitoring tools of the future for Network Planning and Fault Isolation activities, channel utilization has always suffered from an inate limitation that it is a measurement that "pegs" somewhat near 100%, as shown in **Fig. 3** below. (Note: In order to compute the payload channel capacity to be used in our calculation of channel utilization we must discount the raw channel capacity by the bandwidth consumed by FEC overhead, protocol overhead and management messages. This discounting often involves some degree of estimation which introduces a small amount of error into the channel utilization measurement.)

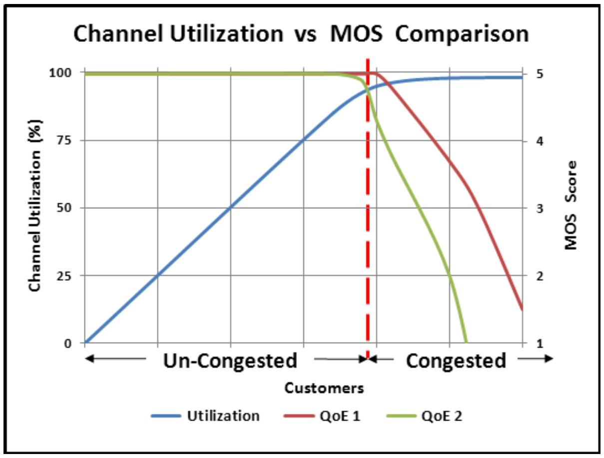

Fig. 3- Channel Utilization & MOS Scores vs. # Customers

Fig. 3 illustrates a hypothetical scenario where every new customer added to the channel adds a fixed amount of "requested load" to the channel. The requested load is the amount of bandwidth that the customer would like to receive to satisfy their desired high-speed data activites. The "composite requested load" on a channel is the sum of the requested loads from all active subscribers, and it should be obvious that (depending on the number of active subscribers at any given time) it can exceed 100% of the channel's capacity. But when the composite requested load exceeds 100% of the channel capacity, it is clear that the actual channel utilization will still be a value that is less than or equal to 100%. This difference between the composite requested load and the actual channel utilization is obviously a good indicator of subscriber QoE degradation, because it represents unsatisfied bandwidth demands.

It should be clear that if a particular customer's requested load is less than their Maximum Sustained Traffic Rate (Tmax) level within their SLA, then they will usually be offered 100% of their requested load if channel congestion (and channel utilization) is low enough. As channel utilization approaches the 95% level, though, congestion starts to develop during short transient periods of time when simultaneous requests for bandwidth from different customers occur. As channel utilization climbs higher than 95%, congestion issues worsen. It is at this point that CMTS congestion control algorithms (such as Weighted Random Early Discard) begin to turn on, delaying and dropping packets from selected Service Flows based on the priority, Tmax, and Tmin settings with each customer's SLA agreement. The dropping and delaying of packets obviously reduces the actual (utilization) load on the channel to be less than the requested load and typically results in degradations in the subscriber QoE levels.

In the past, these CMTS-induced packet drops and packet delays did not cause major QoE problems for subscribers of most network operators, because most network operators kept their channel utilizations well below the 100% level (by, for example, scheduling node splits if average utilization levels exceeded 70%). This had historically been a very viable network planning strategy since the QoE for traditional web traffic types (e.g. Web Browsing and Traditional Streaming Video) degraded very quickly under even small amounts of channel congestion, channel drops, and channel delays.

Today, however, new "compressible" traffic types, such as Adaptive Bit Rate (ABR) Video, have appeared, making it possible for networks to operate reasonably well at what used to be unacceptably high levels of channel congestion. Some of these ABR Video algorithms probe and sense the available bandwidth on the network channels and may even use more bandwidth than required (by switching to HD video resolutions) if extra bandwidth happens to become available. These ABR Video algorithms will also back down to lower resolution video streams if bandwidths are ever found to be lacking due to network congestion.

With continual economic pressures on MSOs, some may be considering ways to minimize their network equipment costs by moving to higher average levels of channel utilization within their DOCSIS networks, relying on the adaptive nature of ABR Video algorithms to benevolently throttle their traffic rates whenever excessive channel congestion occurs. These concepts are feasible, because ABR Video is making up larger and larger percentages of the total network capacity. But can/should MSOs trust the future QoE levels of their subscribers to ABR algorithms that are managed and maintained by third-party content providers on the Internet? Can/will these algorithms throttle their video resolutions (and bandwidths) fast enough to respond to network congestion and ensure that other traffic types (ex: Web-Browsing) are not negatively effected? A study of this topic was conducted by the authors, and it was concluded that Web-Browsing applications multiplexed with ABR Video streams can still suffer in the presence of network congestion. [Clo2] As a result, it seems clear that MSOs who plan to capitalize on the adaptive nature of ABR Video to increase their average channel utilizations must identify techniques to determine the QoE levels of their subscriber applications during the transient periods of time when the channel utilization hits 100%. In particular, they need their QoE Monitoring tools to find ways to guestimate the composite requested load (in addition to the channel utilization) to determine how much bandwidth is being throttled. Obviously, as the composite requested load grows to be much higher than the channel utilization level (which may be pegged at 100%), more bandwidth is being throttled and MOS scores will drop. This is illustrated clearly by the dropping MOS scores on the right-hand side of **Fig. 3**.

Determining the composite requested load on a particular CMTS channel is a challenging problem, because the CMTS congestion control mechanisms that tend to throttle TCP traffic when it exceeds certain thresholds tend to reduce the traffic levels by throttling them at the TCP source. As a result, the actual requested bandwidth levels may not actually "show" themselves to the CMTS or the QoE Monitoring tool. But there are ways that the CMTS may be able to infer and approximate how high the requested load levels actually are. One way is to observe the behavior of queues as the CMTS congestion control mechanisms begin to drop and delay packets. The rate at which CMTS queues were growing can help to give some level of information regarding the current requested load levels. Other techniques could also be utilized. But in general, it should be clear that even an approximate measure of the composite requested load levels on a channel can help the QoE Monitoring tool determine how much QoE degradation is being caused by traffic throttling during periods of high congestion.

<u>MOS Scores Based On High Sampling Rates</u>

The measured bandwidth level or packet delay level or packet loss level for a particular subscriber (or for a particular type of subscriber with a particular SLA) is assumed to be accurate for only a relatively short period of time, because bandwidth levels and delay levels and loss levels tend to fluctuate quite rapidly on DOCSIS channels. Ideally, these traffic metrics would be measured with very fine granularity (on the order of a few seconds or less), but for most QoE Monitoring tools, the ability to collect average channel bandwidth levels is limited by the rate at which statistics can be collected from the devices that are measuring the bandwidths.

This precludes the use of exceptionally high sampling rates.

In the future, bandwidth sampling periods of, say, 1 hour may not be adequate. If a subscriber is not satisfied, they are most worried about the bandwidth levels that are occuring 'now' - not 1 hour ago when the last sample was taken.

Thus, it is possible that future QoE Monitoring tools will be more tightly integrated with the CMTSs and CMs and the other elements that are typically taking the bandwidth measurements.

As an example, future CMTSs and CMs may be required to collect measurements every few seconds and perform some post-processing of the data, sending the QoE Monitoring tools only anomalous results that they observed during their post-processing.

MOS Scores Based On Different Scopes

MOS scores within a QoE Monitoring tool can be created for many different scopes, and they all can have value. As an example, measurements can be taken to determine the traffic attributes (bandwidth, latency, jitter, and loss) at the TCP/UDP session level, at the Service Flow level, at the modem level, at the DOCSIS channel level, or at the DOCSIS Service Group level. Each of these measurements at each of the different scope levels yield different pieces of information.

For example, when monitored at the modem level, these measurements can be used to determine if there is congestion and QoE degradation that is being created by large amounts of traffic entering a single home. This can be useful information that an MSO can use to help determine when a subscriber is a candidate for a Service Tier upgrade recommendation.

When monitored at the DOCSIS channel level or DOCSIS Service Group level, these measurements can be used to determine if there is congestion and QoE degradation that is being created by large amounts of traffic generated within the subscriber's neighborhood. This can be useful information that an MSO can use to help determine when a node-split is required or when extra channels need to be added to augment the capacity of a DOCSIS Service Group.

These two situations (modem congestion vs DOCSIS channel congestion) are quite different from one another and may require different MSO responses, so it may be important for a QoE Monitoring tool to be able to monitor and differentiate between the two situations.

AN EXAMPLE OF A FUTURE QoE MOS SCORES

Future QoE MOS Scores

In the previous sections, we described several ways to augment the typical measurements that are oftentimes utilized within QoE Monitoring tools. In this section, we will give some actual examples of future QoE MOS scores that attempt to make use of bandwidth levels, traffic types, and SLAs.

As suggested in the preceding material, measured per-subscriber metrics can be used as good predictors of QoE. One of the most obvious applications of this within a QoE Monitoring tool is to predict QoE for video viewing subscribers based on the bandwidth that each subscriber is receiving. Another application is to monitor the bandwidth and determine its impact on typical Web Browsing experiences.

From the channel capacity and the channel utilization levels (or from more advanced techniques that utilize active subscriber counts and SLA levels and CMTS scheduling algorithm knowledge), a QoE Monitoring tool can calculate the predicted amount of bandwidth that a particular subscriber (or type of subscriber) is likely receiving at a given instant in time. Once that predicted bandwidth level is calculated, the next obvious question that must be answered is whether that predicted bandwidth level is adequate to offer good QoE levels to the different applications that the subscriber might be utilizing. This question is answered by developing MOS scores for each of the applications, where the MOS scores are potentially a function of the application type, the predicted bandwidth level, and the SLA level for the particular subscriber.

In this section, we will focus on MOS scores for two popular application types (IP Video applications and Web-Browsing applications), but it should be clear that MOS scores can (and will) be developed for the many other types of applications that high-speed data subscribers utilize.

In developing these MOS scores, it was assumed that bandwidth bottlenecks for subscribers are often caused by the narrower bandwidths encountered on the DOCSIS network. This assumption may or may not be correct at all points in time, because it is of course quite possible that there are other bottlenecks in the data path between the source and destination of any Internet exchange. These other bottlenecks could occur in the servers or in the backbone network or in the subscriber's own home network. As a result, the development of MOS scores based solely on measurements taken on the DOCSIS network is, by nature, very subjective, and erroneous conclusions can sometimes be drawn about a user's QoE level. Nevertheless, the authors would argue that there is value in these DOCSIS-oriented MOS scores.

Another potential source of error in these MOS scores could result from low sampling rates (which was discussed above). The bandwidth levels offered to a particular subscriber are expected to vary over time and technological change. As a result, the QoE Monitoring tools must continually be adapted to keep up with the required sampling rates for the traffic types that appear on the networks.

<u>Example Future MOS Scores For IP Video</u>

The authors carried out first-order human factors experiments on several IP Video services. Various bandwidth limits were specified, and the QoE level of the resulting IP Video viewing experience was ascertained. These results are less than statistically perfect, because many more subjects would have been required to carry out a more thorough study. Nevertheless, the results still provide some degree of information about the quality of IP Video viewing with different amounts of bandwidth being offered to the subscriber.

These results quickly revealed that many of the popular video aggregation services (such as Hulu, VUDU, Netflix, etc.) provide different qualities of service (Standard Definition, High Definition, etc.). Some (but not all) of these service can be *adaptive*. In other words, some of them sense the actual bandwidth available on their TCP session propagating over Internet connections, and they decrease or increase the quality of the video delivered based on this bandwidth.

Another important consideration is that some of these services sense the type of consumer device being used and adapt the BW needed accordingly. For example, a tablet with a small screen would be sent a lower resolution (and lower bandwidth) feed than a 'smart TV' with a 56" monitor.

Table 4 contains some experimentally determined, unconstrained bandwidths required for a number of aggregation services and target end-customer devices. These are the bandwidths which the services would (apparently) elect to utilize if there were no bandwidth limitation between the server and the client. For example, end customers who had chosen to use VUDU (a rental movie service, and who had opted and possibly paid extra to receive a super-high resolution video) would require approximately 9.5 Mbps to be fully satisfied with the service. However a viewer watching the same movie in Standard Definition Netflix on a tablet would not benefit from a bandwidth over ~1.5 Mbps. (Note: This table is not exhaustive. As an example, Netflix advertises that they also offer higher-definition streams requiring rates of 3 Mbps, 5 Mbps, and 8 Mbps. However, the experiments performed for this paper were unable to access those higher tier services).

IP Video Application	CPE	Avg Mbps
Netflix	Tablet	1.5
	PC	2.9
	Roku	4.2
Hulu	Tablet	1.3
	PC	1.3
	Roku	2.8
VUDU	SD-PC	2.4
	SD-Tablet	2.4
	SD-Roku	2.3
	HD-Roku	4.7
	HDX-Roku	9.5
You Tube	PC	0.8
	Tablet	2.4

Table 4- Unconstrained Bandwidth Requirements for "Excellent" IP Video Service

Since nearly all of the services had mechanisms for providing *some* flavor of recognizable video down to bandwidths of around 0.8 Mbps (by way of adaptation), viewers would still be able to watch programs down to that bandwidth level, but they may not be very satisfied by the low resolutions. At bandwidth levels below about 0.8 Mbps, most of the IP Video services seemed to break down entirely, resulting in program halts and unsatisfactory QoE levels.

The authors devised a first-order MOS scoring algorithm to provide MSOs with bandwidth vs MOS score guidelines. (Note: It is important to note than any IP Video MOS scoring algorithm will become obsolete soon after it is released, because the IP Video content aggregators are continually changing their ABR algorithms and changing the nature of their content. As a result, IP Video MOS score algorithms must be updated periodically to stay fresh and useful).

As a first-order approximation, the authors would suggest that MSOs utilize the required unconstrained bandwidth levels listed in **Table 4** and associate it with the relatively high MOS score value of 4.8. (Note: Recall that a perfect MOS score of 5.0 corresponds to excellent service). As the predicted bandwidth level of a user falls below the unconstrained bandwidth level within Table 4, then the MOS score will obviously drop. For the purposes of this first-order model, we will assume that the MOS score drops roughly linearly as the predicted bandwidth level drops. We will assume that this linear drop approaches a MOS score of 2 (poor) as the predicted bandwidth level for a subscriber drops to 0.8 Mbps. If the predicted bandwidth level drops below 0.8 Mbps, then the MOS score is assumed to instantly fall to a value of 1 (bad), because the resulting IP Video is no longer deemed to be useable. These MOS scores tended to correlate fairly well with the IP Video viewing experiences of the authors.

From **Table 4**, it becomes clear that the MOS score for a particular subscriber would not only be dependent on the predicted bandwidth level available to the subscriber, but it would also be dependent on what particular IP Video service the subscriber is accessing and the device on which he or she is watching it. This is reflected in **Table 4**.

Example Future MOS Scores For Web-Browsing

The QoE for a Web-Browsing application is predominantly determined by the time between a subscriber's 'click' on a web page's hyperlink and the time when the end-user has the perception that the web page is (for the most part) downloaded and displayed on his or her screen. Since web pages vary widely in size, the QoE for different customers clicking on different URLs is a bit difficult to evaluate. We therefore used published average web page size estimates for our evaluations.

The size of the average webpage has grown over the last few years to roughly 1.3Mbytes, or about 10Mbits. [Inf1] (Note: As with IP Video, web page sizes also vary over time. As a result, QoE Monitoring tools must be periodically upgraded so that they can be trusted to perform their Web-Browsing calculations using appropriate web page sizes).

If we assume that an end subscriber is fully satisfied (i.e.- with a MOS score near 5.0) whenever they perceive the downloading of *an entire* web page of this size to occur within, say, 1 second, then he or she would require a minimum bandwidth of about 10Mbps to achieve this high QoE level.

Shorter web pages or longer web pages would, of course, require proportionately lower or higher bandwidths to achieve the desired one second download time goal.

However, complicating the relationship between bandwidth and Web-Surfing QoE is the fact that, very often, a customer will be quite happy with what he or she sees on the screen BEFORE the entire web page is downloaded. We name this time the 'splash time' and find that (most often) only about 1/2 of the total page's data content is required to create a readable screen, with the remaining bytes for the rest of the screen coming in after the spash time has passed.

We therefore assume that for the average-sized web page of 10 Mbits, the web page size associated with the splash time would be ~5 Mbits. Thus, if a single subscriber is receiving a predicted bandwidth of P, then the average splash time (in seconds) for that subscriber can be calculated by:

$$\text{Average Splash Time} = (5\ \text{Mbits})/P \qquad (3)$$

where the predicted per-subscriber bandwidth P is measured in units of Mbps.

Given these considerations, the authors would recommend that we associate a Splash Time of 1 second to a MOS score of 4.8, and we would also recommend that we associate a Splash Time over 10 seconds (unbearable to nearly all users) to a MOS score of 1.0. Note that these ratings are for average web pages of size 10 Mbits. A resultant table of Mbps vs. MOS would be represented by **Table 5**.

Predicted Subscriber Bandwidth (Mbps)	Average Splash Time (seconds)	MOS Score
>5	<1	4.8
4	1.3	4
3	1.7	4
2	2.6	3
1	5	2
.5	10	2
<.5	>10	1

Table 5- Subscriber Bandwidth vs. Splash Time vs. MOS Score For Web-Browsing

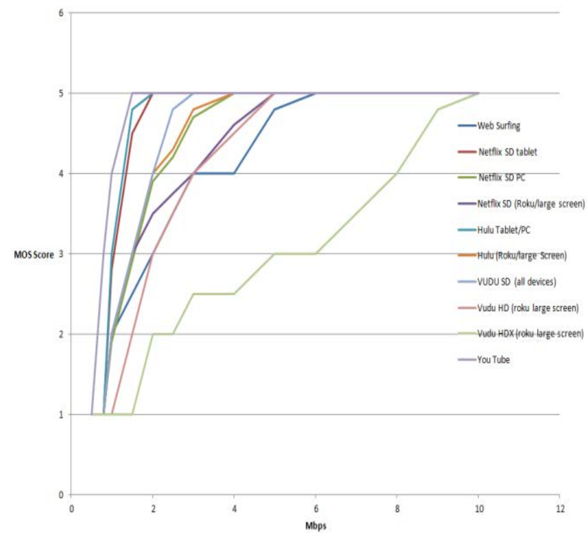

Fig. 4- Bandwidth (Mbps) vs. MOS Score For Different Applications

<u>Combining Future MOS Scores For IP Video And Web-Browsing</u>

For purposes of illustration, **Fig. 4** graphically summarizes the above information relating Predicted Subscriber Bandwidth to MOS scores for various applications. Note, for example, that the light green, rightmost line, represents the most demanding application shown- HDX (Very High Definition) VUDU on a large screen Roku-Connected TV. This application would therefore require a subscriber bandwidth of ~10 Mbps to completely satisfy a customer who had paid for this service. In the middle of the graph, a user involved in Web-Browsing would very likely be exceedingly happy and have a high MOS score with a bandwidth of 6 Mbps. On the left side of the graph, a standard Netflix viewer on a PC with a bandwidth of 2.0 Mbps would be exceedingly happy and have a high MOS score.

A few final points should be made about the application of the plots in **Fig. 4**. If we select a particular bandwidth value (say 4 Mbps), we find that each application can have a very different MOS score. As a result, MOS scores are very sensitive to the application being utilized.

The plots can be used once the bandwidth for a subscriber has been calculated. This calculation can be performed using simple algorithms (ex: total channel bandwidth divided by number of active users) or using complex algorithms (ex: bandwidth calculations based on CMTS scheduling algorithms and SLAs). But regardless of the type of algorithm used, the end result is the creation of a predicted per-subscriber bandwidth level on a particular channel or service group.

The plots can be used to predict the QoE levels for a particular subscriber, in which case the average bandwidth for that subscriber can be calculated by monitoring his or her actual usage levels.

The plots can also be used to predict the QoE levels for generic subscriber types, in which case the average bandwidth for typical users (that share a particular SLA level) would need to be calculated.

Once the average bandwidth (for a particular subscriber or for a generic subscriber type) is calculated, the QoE Monitoring tool can assume that all of the available bandwidth is used by one application, or the QoE Monitoring tool can assume that multiple applications within the susbsriber's home are sharing the available subscriber bandwidth. If the latter is assumed, then more complicated formulae are required to determine how much of the available bandwidth should be allotted to each of the applications that are assumed to be operating within the home.

In any case, there will be a specific amount of bandwidth that is assumed to be available for a particular application. At this point, the MSO or the QoE Monitoring tool can access the plots of **Fig. 4**.

Unless Deep Packet Inspection tools are being used in conjunction with the QoE Monitoring tool, it is unlikely that the QoE Monitoring tool will actually be cognizant of the traffic types that are actually passing to the subscriber at any point in time. As a result, the plots in **Fig. 4** can be utilized to describe the QoE levels for the subscriber IF they were currently accessing different traffic types (whether they are actually accessing those traffic types or not). Thus, the resulting MOS scores that are pulled out of the plots in **Fig. 4** are actually the MOS scores for hypothetical subscribers assumed to be using that service - they are NOT actual subscriber MOS scores. It is important to note this distinction.

If an MSO does not know exactly which application a subscriber is using at any given time, an obvious question comes to mind. How can the MSO predict the QoE levels of a subscriber without that knowledge? One way to do it is to create the concept of a "typical subscriber model." An MSO can construct this model using knowledge about the mixes of applications that their aggregate set of subscribers use on the Internet. For example, if an MSO knows that (on average) the mix of traffic types is given by 54% IP Video traffic, 18% Peer-to-Peer traffic, 11% Web-Browsing, 17% other traffic, the MSO can assume (for modeling purposes) that a "typical subscriber" might be temporally transitioning between the different traffic types such that his or her average bandwidth mix equals the percentages above. For the 54% IP Video Traffic mix listed above, the MSO could further try to determine the mixes of different types of IP Video traffic on the network (ex: Netflix, VUDU, etc.), and the MSO could then assume that the "typical subscriber" is temporally transitioning between video streams from those different IP Video traffic types whenever they are supposed to be viewing IP Video. As a result, a blended MOS score can be created by creating a weighted average of all of the MOS scores for all of the heavily used applications. This blended average can be useful in ascertaining whether a "typical subscriber" is currently experiencing high QoE levels or not. If more detailed information is desired, then the MSO can dive into the details of each individual MOS score that contributed to the weighted average.

EXAMPLE OF FUTURE QoE MONITORING TOOL DISPLAYS

This section discusses an implementation of a QoE Monitoring tool that is possible to achieve with today's CMTS technology. We base this on the observation that every existing CMTS must currently incorporate some mechanism for resolving downstream channel congestion. We further observe that (regardless of the congestion resolution mechanism used) we should be able to model, to a reasonable degree, the current state of that mechanism as a mapping of the DOCSIS SLA parameters (Priority, Tmax, Tmin, etc.) from their SLA-specified values to a set of "Effective Values" that will limit the aggregate downstream data rate to fit into the available downstream bandwidth. In order for this process to be "fair", rather than random and arbitrary, we believe that most CMTS systems must currently maintain, in some fashion, a set of "state" information that would reveal this mapping for the current level of channel congestion.

While the volume and complexity of the information required to communicate this mapping may vary widely among CMTS vendors, we have shown, by the example described here, that in some cases this mapping can be summarized and communicated in a surprisingly small amount of data for each downstream channel. In the experiment described here we were able collect this congestion state information for all CMTS downstream channels at intervals as short as 15 seconds – which allows us to collect congestion-state information with a high degree of time resolution and without burdening the CMTS with a large volume of I/O.

Once we have this capability of mapping SLA-specified DOCSIS parameters (Priority, Tmax, Tmin, Tpeak, maxBurst, etc) into Average Bandwidth per Subscriber Service Flow values based on the current level of channel congestion, it is then a relatively simple matter to "predict" the length of time required to transfer a data block of any given size and, in turn, the resulting equivalent data rate and the resulting response time for various Web applications.

If we now select some data block size that is "typical" for a given type of internet service (e.g. a typical web page size for web browsing, or a 2-second block of video data) we can predict, with reasonable accuracy, both the response time and effective data rate that a customer would experience in requesting that data block (as a function of that customer's SLA parameters).

Finally, a table mapping the effective data rate or response time for that hypothetical internet service to a MOS Score (based on the subjective judgement of representative users) can yield a useful measure of a customer's expected QoE.

SLA	Surf	NFlxHD	NFlxSD	VuDuHD	Speed	YouTube	ABR_7	ABR_3
Gold	5	5	5	5	4.86	5	5	5
Silver	4.59	4.59	4.59	4.59	4.59	4.59	4.59	4.59
Bronze	4.59	4.59	4.59	3.92	4.59	4.59	1.02	4.59

Fig. 5- MOS Score Matrix Display

The complete QoE status for a channel at any given point in time can be displayed in the form of a matrix as shown in **Fig. 5**. This display is able to simultaneously display a MOS Score for every combination of SLA (Gold, Silver or Bronze) and Internet Service Type. (The service types corresponding to the eight column headings shown in this example are, respectively: Web Surfing, SD Netflix, HD Netflix, HD VuDu, Speed Test, YouTube, 7 Mbps ABR Video, and 3 Mbps ABR Video).

It is important to remember that a customer's quality of experience is affected not only by network congestion but also by their choice of SLA subscription. In the example in **Fig. 5** the customer's anticipated modest satisfaction (3.92) with VuDuHD and complete dissatisfaction (1.02) with any attempt to view 7 Mbps ABR Video is a consequence of the a Tmax value of 5 Mbps in their SLA.

While the matrix format shown above can accommodate a large number of SLA and Service Type definitions it lacks a time dimension that would permit the display of chronological trends for individual values. A chronological display on the other hand, as shown in **Fig. 6**, can display a wealth of time dependent detail, but becomes confusing if we attempt to include too many SLA/Service Type combinations. (This example has only included five of the eight Service Types.)

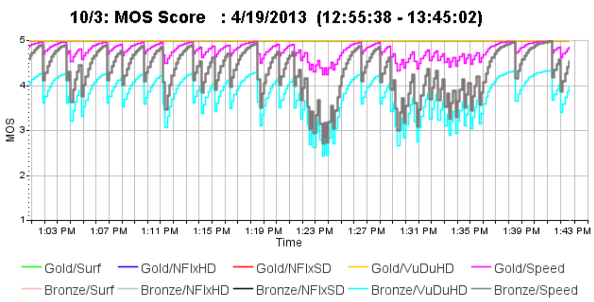

Fig. 6- MOS Score Chronological Display

The very short sampling intervals (as low as 10 seconds) possible with this approach can provide a high degree of chronological resolution without massive data collection overhead – thus capturing congestion detail that would be completely lost with longer sampling intervals of 15 minues or an hour.

CONCLUSIONS AND FUTURE POSSIBILITIES

This paper has attempted to answer several important questions about next-generation DOCSIS QoE Monitoring tools:

1) Why should we suddenly be interested in a family of MOS scores and metrics that we have comfortably ignored for so long?
2) What is it about our DOCSIS networks that has changed to require the use of more sophisticated MOS scores and metrics?

Both of these questions have the same answer. DOCSIS networks have historically used only a single primary metric (channel utilization) when attempting to determine QoE performance – and that metric is quickly losing its efficacy due to several facts:

1) More bandwidth capacity is being offered within the SLAs of DOCSIS subscribers
2) More applications with very different characteristics and a wider range of bandwidth needs are now sharing DOCSIS channels
3) More bandwidth is being used by each subscriber application
4) Each application is sensitive to different QoE metrics (bandwidth, delay, jitter, packet loss, etc.)
5) Each application has different acceptability thresholds defining good QoE levels
6) Benevolent applications (like ABR IP Video) that yield bandwidth to others are sharing channels with greedy applications (like Peer-to-Peer)
7) ABR IP Video is a novel traffic type that is “sponge-like”- growing to use more bandwidth when the bandwidth is available and shrinking to use less bandwidth when the bandwidth is not available
8) Single subscribers can now have a strong negative impact on all of the other subscribers sharing a channel

The authors believe that QoE Monitoring tools and QoE metrics will need to undergo a rapid change to help fill the instrumentation void caused by basic inadequacies of simple channel utilization measurements from the past. In the future, MSOs are likely to need multiple QoE metrics for their DOCSIS high-speed data service tier (instead of the single metric of channel utilization that is predominantly utilized today).

Because DOCSIS networks are carefully designed to always provide customers with the highest quality of service for which they have subscribed (unless it is physically impossible to do so), one might expect any measure of QoE to indicate full customer satisfaction unless global channel congestion levels or inadequate DOCSIS Tmax levels are limiting the subscriber's bandwidth to a level less than desired. If it is the former (channel congestion) that is causing the QoE problem, then the MSO may need to make changes to upgrade their network. If it is the latter (inadequate Tmax levels) that are causing the QoE problem, then the subscriber may need to make changes to upgrade their SLA level.

As a result, it is necessary for QoE Monitoring tools to make use of channel utilization as a metric to indicate whether the problem lies in the MSO's camp or the subscriber's camp. But monitoring of other metrics (either direct or indirect) can also be used to help the MSO know exactly which types of applications may be experiencing problems at any point in time. When coupled with more detailed information about subscriber SLA levels and CMTS scheduling algorithms, these tools can become quite intelligent as they predict the QoE levels for subscribers.

This paper has shown that many different improvements and augmentations can be added to future QoE Monitoring tools. In addition, the paper has shown how MOS scores of the future can take advantage of the many improvements that are listed. Finally, the paper showed an example of a QoE Monitoring tool and illustrated some useful display models.

Several key ideas were outlined within this paper. First and foremost, it was shown that knowledge of traffic types, subscriber SLA levels, and traffic metrics (such as bandwidth, delay, jitter, and packet loss) can help QoE Monitoring tools create much more specific and accurate predictions about subscriber QoE levels. In addition, it was shown that the ability of a QoE Monitoring tool to predict the QoE performance levels for different traffic types can be greatly improved if the QoE Monitoring tool is made cognizant of the particular scheduling algorithms used by the CMTS. With this knowledge, the QoE Monitoring tool can make more informed decisions about how the traffic will be handled by the CMTS during periods of congestion and non-congestion.

The fact that the QoE Monitoring tool is improved by receiving information from (and about) the CMTS leads the authors to wonder whether the opposite is also true. If information from the QoE Monitoring tool could be rapidly coupled back into the CMTS, could the CMTS make good use of that information? Is it possible that future CMTSs could dynamically modify their scheduling algorithms to try to assist a particular traffic type that might be experiencing low QoE levels at a particular instance in time? Or is it possible that load-balancing algorithms inside of CMTSs could make good use of the outputs from the QoE Monitoring tool to re-balance CMs attached to the channels on a CMTS? Even further, is it possible that future CCAP elements could initiate dynamic QAM sharing functions to modify the number of EQAM channels and DOCSIS channels in response to low DOCSIS MOS scores from the QoE Monitoring tool? Only time will tell if these interesting concepts might eventually find their way into real-world applications as QoE applications continue to evolve.

REFERENCES

[Clo1] T. Cloonan and S. Chernak, "HFC Spectrum Planning in the Era of the HSD Bandwidth Explosion," Proceedings, SCTE Cable Tec Expo '12, October, 2012.

[Clo2] T. Cloonan and J. Allen, "Competitive Analysis of Adaptive Video Streaming Implementations," Proceedings, SCTE Cable Tec Expo '11, November, 2011.

[Inf1] http://www.infodocket.com/2013/03/06/statistics-the-average-size-of-a-web-page-is-now-1-3-mb-up-33-in-last-year/

Analog Reclamation via IP DTA

Ron Shanks, Sangeeta Ramakrishnan
Cisco Systems

Abstract

Analog Reclamation is at the forefront of most Cable Operators minds today. In a 750MHz cable plant it is likely that half or more of the available spectrum is allocated to basic or extended analog programming. Obviously this spectrum is attractive to reclaim in order to add additional tiers of programming. Internet bandwidth usage is also growing rapidly, with peak hour usage growing at an annual rate of 40-50%. Analog reclamation could provide operators with the spectrum needed to meet consumer demands. Until this point one of the only options available for analog reclamation was to use a DTA, also known as a Digital Transport Adapter.

Traditional DTAs solve the analog reclamation problem by allowing an operator to move traditional analog subscribers to QAM delivered video with a very inexpensive device. Unfortunately the DTA device is limited in that there is no two-way communication capability with the network. For the subscriber this translates into no advanced Guides, no On Demand Capability and no way to consume OTT content.

This paper will describe an alternative to the traditional DTA called the IP DTA. This alternative leverages the operator's investment in DOCSIS 3 by using Bonded DOCSIS channels as the transport for the legacy analog tiers. The concept takes advantages of cost reductions in newer IP Set Top boxes coupled with cost reductions gained by using Downloadable Conditional Access (DCAS) for security. The paper will outline the network model and will demonstrate the ability to support the entire analog tier being carried on as few as 4 EIA channels. The paper will also highlight financial benefits and operator savings that can be realized by moving straight to IP versus going through a traditional DTA deployment.

Introduction

Today's Cable Operators are challenged with a growing demand from their subscribers for more High Definition programming and increased bandwidth for Internet access.

The fundamental problem is a lack of spectrum on the HFC plant. A typical 750 MHz cable plant supports approximately 116 EIA channels. The largest single allocation is the spectrum allocated for traditional analog programming. In many cases greater than 50% of a carrier's spectrum is allocated to carrying analog programming. Due to the size, it is this portion of the spectrum that is the most appealing to convert for other purposes such as increased DOCSIS capacity for IP Video and High Speed Data. Unfortunately most carriers still have a significant percentage of their subscribers who are "Analog Only" or have televisions that use the analog service.

The top of mind question for many operators is how to reclaim and repurpose this analog spectrum while continuing to provide services to the traditional analog subscriber base? The following sections will outline the different technology alternatives available to address this challenge, detailing the pros and cons. The paper will also introduce the IP DTA as a viable solution based on recent technological innovations.

Figure 1 Typical spectral allocation

Potential Solutions

Analog Reclamation using Traditional DTAs

One option to address the bandwidth constraint is to connect a Set Top Box to the TV and get the programming from the Digital QAM. This approach works because the programming is already replicated on the QAM. The Traditional Digital Terminal Adaptor (DTA) was created to serve as an inexpensive one way digital QAM fed CPE that does not have to support more expensive Content Protection capability. Because the analog programming is not encrypted, there is not a requirement to encrypt this content when it's carried digitally. While the Traditional DTA allows for encryption, the encryption technology is usually a much lighter weight technology than the traditional Conditional Access technology that is used to protect the higher tiers of programming delivered via Digital QAM.

Traditional DTA Advantages

Analog reclamation using DTAs is not a new concept and the technology is relatively mature. The prices for DTAs themselves are relatively inexpensive. Also deployment of DTAs requires minimal to no changes to the transport network. In almost all cases DTAs can be "self-installed" by the subscriber and do not require a truck roll to be deployed.

Traditional DTA Disadvantages

One of the biggest disadvantages of the current Traditional DTA is it is a one way device. This significantly limits its capabilities going forward and potentially pushes the DTA into early obsolescence. For example, the Traditional DTA is not able to receive a next generation two-way programming guide. The lack of such two-way capability also means that the Traditional DTA cannot be used to consume On Demand services which could result in missed opportunities for additional revenue for the operator. Because the Traditional DTA does not have the capability to support high security Conditional Access (CAS) it cannot be used to receive higher programming tiers that require CAS, again limiting the revenue potential for the operator. Also the lack of capability for a Traditional DTA to support two-way OTT (Over the Top) video services severely limits the platform. Many operators are beginning to view Traditional DTAs as a short term fix for the analog reclamation problem. Many operators view purchasing Traditional DTA hardware as throw away capital because the Traditional DTA does not have the ability to support many of the next gen video features and ultimately limits up-sell opportunities.

1 GHz Expansion

Although plant expansion is not typically considered a reclamation strategy, it is a valid approach to increasing available spectrum. Additional spectrum that can then be used to expand programming and High Speed Data services. Expansions to 870 MHz and 1 GHz have been performed by many North American operators. A 750 MHz to 1 GHz plant expansion provides approximately 40 more EIA channels to be used for other purposes.

1 GHz Advantages

Spectrum expansion provides the opportunity for the additional capacity to enable more programming and higher speed data services. The technology to enable this extension is reasonably mature. This strategy also allows the operator to expand capacity without having to change the analog programming tier, such that an analog reclamation strategy becomes additive.

1 GHz Disadvantages

The major drawback to the spectrum expansion strategy is cost. A migration to a 1 GHz plant is costly per subscriber, especially if coordinated with a move to smaller node sizes. Also, most legacy CPEs – (QAM Set Top Boxes and Cable Modems) cannot take advantage of the new spectrum. This means that operators will need to introduce new CPE hardware (at additional cost) to take full advantage of the upgrade.

Switched Digital Video (SDV)

Traditional QAM based Switched Digital Video (SDV) technologies take advantage of low concurrency for long tail programming, allowing less popular content to be switched on and off the HFC plant. Only channels that are currently being viewed take up capacity. This allows an operator to offer significantly more "low viewership" programs. It is possible to deploy SDV to reclaim spectrum that could be used to increase high speed data capacity, but most of the time, SDV is deployed purely to extend the programming lineup.

SDV Advantages

The advantage to SDV is that it provides a more efficient way to support more programming choices on the same spectrum by allowing less popular programming to only traverse the HFC network when a customer is viewing it. This approach allows the operator to offer more total programming choices. SDV technology is relatively mature and well understood.

SDV Disadvantages

While the technology for SDV is mature it is also very complex and expensive. It is worth noting that SDV technology is also proprietary from vendor to vendor. While the switching function provides a significant benefit because it targets low concurrency content it typically is not bundled with other network efficiency strategies such as newer and more efficient video compression. SDV also requires an expensive CPE that requires a Cable Card for content protection in the US market.

IP Digital Terminal Adaptor (IP DTA)

It is widely accepted that a migration to an all IP Video infrastructure is the desired endpoint for most operators. The challenge lies in the migration to IP while continuing to leverage a significant portion of the legacy investment in network and CPE. One strategy is to employ IP Video over a DOCSIS network. This is not a new concept, but up until recently hasn't qualified as a valid "Leap Frog" technology.

The concept is relatively simple. It requires creating additional virtual IP paths to the home using the bonding technology provided by the migration to DOCSIS 3.0. Using traditional IP Multicast Video technology this approach takes advantage of the concurrency gain for popular or "short tail" content. It also leverages newer compression technologies such as H.264 to further increase efficiency. In addition, IP Video over DOCSIS takes advantage of newer, more powerful and cost effective CPE such as DOCSIS 3.0 Cable Modems, Gateways and IP Set Top Boxes. Finally, this approach takes advantage of new Downloadable Conditional Access (DCAS) technology which eliminates the need for proprietary hardware or a traditional Cable Card.

The focus of this paper is centered on exploring a primary use case for leveraging current Video over DOCSIS technology as an alternative to deploying Traditional DTAs for analog reclamation. The premise is to retask reclaimed analog spectrum into DOCSIS capacity that will be used to simulcast the previous "Analog Tier" as IP Video over DOCSIS. The operator will then deploy an IP STB (i.e. IP DTA) to terminate this programming in the home.

Figure 2 Spectral allocation with IP DTA

IP DTA Advantages

There are many advantages to adopting IP over DOCSIS technology as an Analog Reclamation strategy. While the IP DTA is more expensive than the Traditional DTA it is much more "future proof" and is not considered throw away capital. The IP DTA is a two-way device that allows an operator to significantly enhance the customer experience compared to a Traditional DTA. At a minimum, IP connectivity allows operators to deploy advanced two-way programming guides and enable two-way services such as VoD. It also enables the operator to open up a much wider portfolio of programming and services that haven't been available on traditional QAM Set Tops, such as OTT content or content stored and delivered by other devices in the home. Video over DOCSIS can also be deployed with dynamic IP multicast so that only the content being viewed by a subscriber is put onto the HFC plant. This approach maximizes the efficiency of the simulcasted programming.

Most believe that Traditional DTAs have a limited life span and that the total CPE cost over time is less if an operator bypasses Traditional DTA deployments and moves straight to an IP DTA.

The migration to an IP infrastructure not only enables analog reclamation, it also introduces possibilities for future growth. As the operator expands the IP migration beyond just the Analog Reclamation use case the advantages significantly increase. The ability to support very secure content protection using downloadable CAS allows the operator to offer higher tiers of programming and an opportunity for increased ARPU. The two-way capability of the service also introduces On Demand Services, again allowing for an increase in ARPU. Going forward, migrating to an IP headend will also allow for a faster migration to headend based services such as Time Shift TV and Network/Cloud DVR. Centralization of these functions supports the migration of programming storage from the home to the network, ultimately providing a significant long term CPE savings due to reduced CPE storage needs and costs. In addition the IP headend allows the provider to provide IP based Video Services directly to additional devices. In addition to supporting the IP DTA, other common devices could be supported such as Tablets, PCs, Game Consoles, Connected/Smart TVs and IP Dongles. Over time this migration further reduces the home CPE costs as BYOD (Bring Your Own Device) technology replaces STBs. Not only does a migration to an IP Headend support the introduction of additional devices but it also enables advanced services such as

Advanced/Targeted advertising, further increasing additional revenue opportunities.

IP DTA Disadvantages

The primary disadvantage to using the IP DTA in conjunction with IP Video over DOCSIS as an Analog Reclamation strategy is the upfront cost. IP Set Tops are more expensive than a traditional QAM DTA in an "apples to apples" comparison. It is only when the deployment of the IP DTA is seen as a first step towards an overall migration to IP that the benefits become clear. Many operators are in the process of deploying IP Video services to unmanaged devices such as Tablets, PCs, and mobile devices and are already putting in place the headend IP technologies necessary to support IP Video services.

IP DTA Technology Enablers

In the last several years a number of technological developments have occurred which enable the use of the IP DTA as a viable option for analog reclamation.

1) Channel Bonding

In DOCSIS 3.0 [1], channel bonding was introduced in both the Downstream and Upstream directions. With this capability a number of channels can be bonded together to create a high bandwidth pipe over the HFC. Modems with this capability can simultaneously receive 4 or more downstream channels. This is a critical capability for the delivery of IP video due to the high bandwidth nature of video. With channel bonding not only is the pipe wider, but it is more efficient due to the statistical multiplexing gains achieved when aggregating a large number of streams together.

2) Multicast Support

DOCSIS 3.0 also includes a significant enhancement in the multicast support in DOCSIS. It includes support for Source Specific Multicast, and the ability to apply Quality of Service (QoS) to multicast flows. Video remains a one-to-many application where a large number of subscribers are watching the same set of video streams simultaneously. Such concurrence makes IP multicast suited for the delivery of video over cable networks. This rings true for the Analog Tier as well, especially because time-shifted content viewing is not typically available for the Analog Tier. By using IP multicast, the efficient delivery of the Analog Tier content over DOCSIS can be achieved.

A number of options are available for the delivery of multicast. Those include static multicast, dynamic multicast and the use of narrowcast or broadcast channels.

a. Static Multicast

The streams corresponding to the analog tier could be statically forwarded on the DOCSIS network using IP Multicast. Static multicast does not mean that the streams are propagated in the home network statically; rather they are carried over the HFC statically. Static multicast has the advantage that it is easy to provision because the bandwidth required for the video streams is known and is independent of the viewing patterns and popularity of these video streams in various fiber nodes. However the disadvantage of static multicast is that bandwidth is consumed irrespective of whether a particular stream is being requested by any user or not.

b. Dynamic Multicast

The video streams can be delivered via dynamic multicast whereby the streams are forwarded on a segment of the network only when a user requests it. This approach still provides multicast gains in that multiple users (on the same segment of the network) when viewing the same content receive the same replication. Another advantage to this approach is that if no users on a segment request a stream, then it is not forwarded on that segment thereby saving bandwidth that can be used for other services. Dynamic multicast can more effectively save bandwidth when service groups are small, and the lineup includes less popular channels. Dynamic multicast generally trades off spectrum savings for higher costs (due to narrowcast DOCSIS channels) when compared to static multicast.

3) RF Spanning

RF Spanning refers to the approach where a set of DOCSIS downstream channels are split across a number of Fiber Nodes. With RF Spanning a very large virtual service group can be created when a set of DOCSIS channels are split across all (or many) Fiber Nodes served by a single CMTS.

By bonding a few RF spanned DOCSIS channels and forwarding static multicast streams over this Bonding Group, an architecture similar to the broadcast architecture for traditional QAM video can be achieved. This is very cost effective because now the cost of delivering these IP video streams is now amortized across a large number of subscribers, perhaps across all the subscribers served by a single CMTS (Cable Modem Termination System).

Overall a very efficient delivery mechanism can be built by combining all of the features of RF Spanning as described above. (A more detailed description of many of the DOCSIS features that are critical for IP video delivery can be found in [2]). The most popular streams from the Analog Tier can be delivered as static multicast over the RF spanned channels, while the less popular analog channels can be delivered via dynamic multicast over the narrowcast DOCSIS channels. Both the narrowcast and RF spanned channels themselves can be bonded together to get the statistical multiplexing gains.

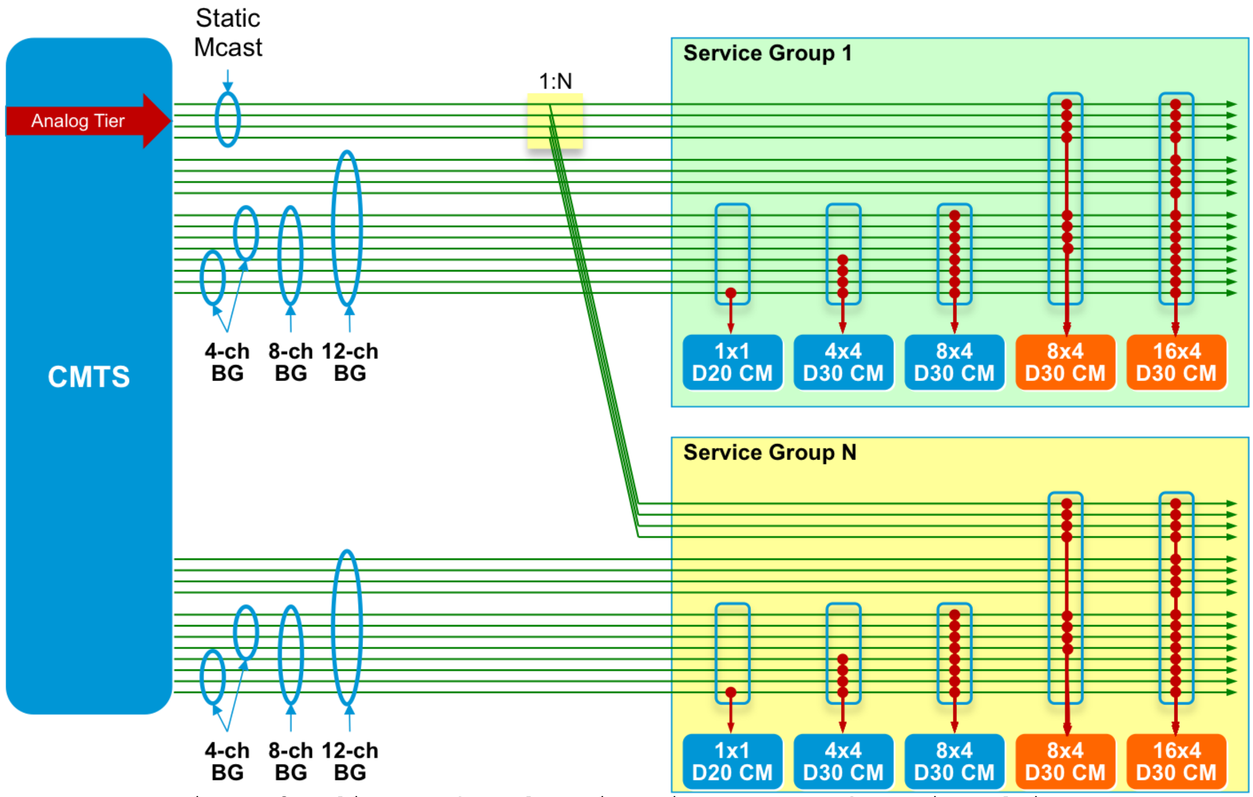

Figure 3 Delivery of Analog Tier via RF spanned static multicast

4) High Capacity Cable Modem technology

Video over DOCSIS had not been widely deployed in the past because high scalability Cable Modem technology was not available. Current 8x4 modem technology quickly becomes insufficient in a Video over DOCSIS model as higher tiers of programming and On Demand Services are added but could be used for the basic Analog Reclamation use case. Manufacturers are bringing Cable Modems to the market that support up to 24 or 32 downstream channels and up to 8 upstream channels. Operators are already evaluating and deploying this technology to offer enhanced High Speed Data offerings and this deployment of higher capacity devices will simultaneously enable managed IP Video services.

5) IP STB Technology

In addition to DOCSIS CPE technology, IP STB technology has also seen a number of recent enhancements. The next generation SOCs (Systems on a Chip) currently available from the major silicon vendors have enabled STBs with significant embedded processing capability. The addition of the new SOC allows these IP STBs to run very advanced middleware software that can process traditional linear video delivery and next generation cloud enabled services simultaneously. In addition the IP STBs also support the current and next generation 3D Electronic Programming Guides (EPGs). Other enhancements include support for advanced home networking technologies such as MoCA (Multimedia over Coax Alliance) 2.0 allowing customers to use existing home Coax to interconnect the devices. IP STBs are also beginning to support advanced Wi-Fi technologies such as 802.11n and the new

802.11ac standard. These advanced technologies also provide the potential for customer “self install” capability therefore reducing the cost of deployment.

6) Downloadable Conditional Access (DCAS) and Digital Rights Management (DRM)

Perhaps one of the most significant developments has been in the area of Content Protection or Conditional Access. Traditional Conditional Access technologies have historically relied on proprietary hardware based solutions and specifically, in the US market, has required the security to be separable which is achieved by utilizing a Cable Card in every STB. Vendors have begun publishing open standards for the hardware component of Conditional Access and silicon vendors have adopted them, making Downloadable Conditional Access a possible alternative to Cable Cards. This allows for a very strong Conditional Access system to be deployed on a generic hardware platform without the traditional Cable Card. In the US market there are examples of FCC waivers that have been granted to operators allowing for Downloadable Conditional Access to be deployed instead of Cable Cards.

For Unicast services such as Video on Demand, “long tail” linear or Time Shifted/Cloud DVR services, DRM (Digital Rights Management) technology can be applied for content protection. DRM is much better suited for Content Protection of Video services delivered via Unicast technologies. The next generation IP STBs have the capability to terminate both types of content protection mechanisms.

7) IP Error Recovery and Rapid Channel Change

Technology initially developed to support IPTV deployments in the Telco/DSL deployment markets has found a new life as well. Edge based servers can enhance the user experience and robustness of linear IP multicast delivery. By sending Unicast bursts to individual STBs, these systems allow for recovery of lost and corrupted linear IP Multicast packets. This same system can also be used to send Unicast bursts to a client enabling almost instantaneous Channel Change times. While not mandatory for an IP DTA deployment, these systems do allow for a mechanism to enhance a viewer’s Quality of Experience.

8) Network Innovation

In the past, IP video migration had been harder to justify given the costs of the DOCSIS network and the scale of the CMTS products available. However in the last several years products with increasing density, scale and lower costs per DOCSIS channel have become available. The Converged Cable Access Platform (CCAP) [3] is in fact recognition by the industry of the need for high density, converged platforms that can aid in the IP video migration. The eventual goal with CCAP is to have all RF channels for a fiber node be stacked on a single spigot on the CCAP platform. The ever increasing density of the CMTS and CCAP products eliminate a major bottleneck in the IP video migration path.

IP DTA Functional Overview

Below is a diagram that outlines the high level functions for an IP DTA deployment that utilizes IP video over DOCSIS to support linear content.

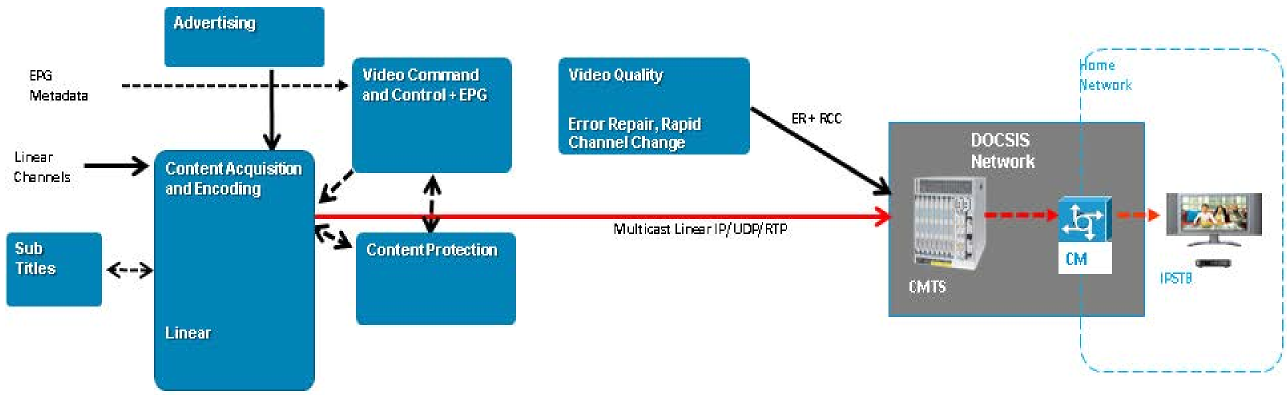

Figure 4 IP DTA Functional Diagram

The basis for linear content delivery is MPEG-2 Transport Streams delivered via IP Multicast. It is expected that most deployments would use an advanced compression technique such as H.264 as it is mature and supported widely on both the encoders and CPE. It is expected that a standard deployment would support Zoned Advertising, Sub Titles, Emergency Alerts and Closed Captioning, thereby matching current capabilities provided via Digital QAM delivery. In addition to these basic features, optional Error Recovery and Rapid Channel Change technology is available from several vendors in the market but is likely to be proprietary.

The in-home network requires a DOCSIS 3.0 Modem/Gateway to terminate the DOCSIS channels. DOCSIS 3.0 is required for most IP DTA deployment scenarios, especially for those that would eventually include offering higher tiers of programming and On Demand content. MoCA technology can be used as a wired networking technology inside the home in addition to wired Ethernet. It is expected that MoCA would be preferable as most subscribers already have coax deployed in the home. An option for Wi-Fi connectivity will also be offered to enable wireless deployment of the IP STB and eventual wireless support for other devices such as tablets, mobile phones, PCs, etc... It is expected that 802.11n technology moving to 802.11ac will be used to support Wi-Fi enabled IP DTAs.

Network Capacity Requirements

One of the important motivations for analog reclamation is to reduce the spectrum inefficiency that is a characteristic of analog video. Hence it is important to understand the spectrum requirements for the IP DTA proposal.

Given that this would require new DTAs to be deployed it is very likely that the DTAs would be H.264 capable, and hence all streams would be delivered via H.264 which is about twice as bandwidth efficient as the MPEG2 codec. Typically Standard Definition video can be encoded with good quality at about 2 to 2.5 Mbps. High Definition video will likely require about 6.5-7 Mbps with constant bit rate encoding.

The bitrates mentioned above correspond to constant bit rate (CBR) encoding, which requires more bandwidth than variable bit rate (VBR) encoding. The use of VBR streams generally provides significant bandwidth savings but comes with a corresponding challenge in managing a multitude of VBR streams in a fixed bandwidth pipe. As shown in [4] when delivering video over IP, VBR encoding can be used to reduce bandwidth consumption by about a third compared to using CBR encoding. The statistical gains from multiplexing a larger number of streams in a fat pipe (enabled by DOCSIS 3.0) make these bandwidth savings a reality.

The spectrum required for IP DTAs depends not only on the bitrate of the streams (SD vs. HD, CBR vs. VBR) but also on the architecture used for the delivery of these streams.

Static Multicast with RF Spanning

Most operators have an average of 65 analog video channels on their HFC (Hybrid Fiber Coax) plant today. Assuming that all of the video channels are carried over IP for IP DTAs, this scenario would require that the HFC plant carry 65 IP video streams. Given that analog video by nature is not high quality video, it may be sufficient to deliver these streams in the Standard definition format. If these streams are encoded in H.264 and in Standard Definition format, then they would require approximately 2.25 Mbps per stream, which translates to 146 Mbps for the entire lineup of 65 streams. The bandwidth required for statically forwarding all of the streams fits well within a 4 channel Bonding Group, assuming 256 QAM 6MHz channels.

Therefore the entire Analog Tier lineup can be transmitted in Standard Definition over IP using 4 DOCSIS downstreams on a CMTS to feed all of the Fiber Nodes on the CMTS. This is very appealing because now the spectrum requirements for the Analog Tier can be reduced from 65 RF channels to just 4 RF channels.

As discussed earlier when statically multicasting streams, the operator could RF span these DOCSIS downstreams that carry the video to many or all Fiber Nodes served by the CMTS. Hence with RF Spanning only 4 DOCSIS channels need to be provisioned for an entire CMTS. This not only reduces the cost for the operator but also helps preserve CMTS capacity for other narrowcast services.

If IP DTAs were being pursued ahead of the overall IP video migration, then it may be efficient to stream these streams in Standard Definition format, rather than High Definition format. However if IP DTAs are being launched along with an overall IP video migration initiative then it is worth considering streaming all of the Analog Tier content in a High Definition format. The reason for that approach is that many of the channels will already be sent in High Definition format, so instead of carrying both formats of content, it may be more efficient to simply carry the High Definition format and have the IP DTA downscale it to Standard Definition as needed.

Today often times the same content is carried 3 times – once as part of the analog tier, second as Standard Definition digital video, and third as High Definition video. We are proposing that when deploying the IP DTA along with an overall IP video migration for the entire lineup, it may be prudent instead to carry a single copy of the content in the High Definition format. Of course the Standard definition and High Definition digital video channels delivered via QAM video will still need to be carried to be able to support the vast number of digital STBs already deployed.

The following table shows the number of RF channels required if all 65 analog channels are delivered as a combination of all Standard Definition, an equal mix of Standard Definition and High Definition, and all High Definition.

Format	RF channels required (CBR)	RF channels required (VBR)
All SD	4	3
Half SD and Half HD	8	6
All HD	12	8

Table 1 RF channels required

As can be seen from the table above, the capacity required for delivering the Analog Tier via IP multicast is well within the scale of today's CMTS and Cable Modem products.

Combination of Static and Dynamic Multicast

Another option is to carry the lineup as a combination of static multicast and dynamic multicast. This would be more spectrum efficient than the "all static multicast" approach if several of the streams in the Analog Tier are in fact long-tail content and are not being viewed in all Service Groups at any given time.

Dynamic multicast has similar gains as Switched Digital Video since streams are not forwarded to Service Groups if no viewer is actively viewing those streams.

Depending on what percentage of the Analog Tier content is long-tail versus short-tail and what kind of over-subscription gains can be achieved on the long-tail content, the number of RF channels required for the service will vary, but will be less than the number of RF channels required in the static multicast case.

Summary

This paper has documented a fundamental problem facing most Cable Operators today. In most scenarios, Cable Operators have exhausted the HFC spectrum that is required to offer additional programming and increase High Speed Data capacity. One alternative is to expand the HFC spectrum out to 1 GHz and provide "new" spectrum. This approach is burdened with high infrastructure cost combined with a need for new CPE that would be required to use this "new" spectrum. Another option for addressing the bandwidth constraint is to deploy Switched Digital Video technology however this option is not likely to provide the significant spectrum re-allocation necessary for increased High Speed Data.

Two strategies for Analog Reclamation include the Traditional DTA and the IP DTA. The traditional DTA is a viable and mature technology for reclamation however the Traditional DTA device is one-way and does not support advanced programming guides or next generation two-way services. The Traditional DTA is considered a short term solution to the problem until operators can eventually migrate to IP delivery.

The IP DTA approach provides the benefits of Analog Reclamation coupled with a more future proof CPE instantiation. This approach begins with an Analog Reclamation strategy which allows for the deployment of a powerful and inexpensive IP STB (the IP DTA) which receives its linear programming via IP Multicast Video over DOCSIS. While the initial cost of this solution seems to be more expensive than a traditional DTA deployment, this is more than offset in the long term. Long term CPE costs are lowered by not having to deploy traditional DTAs in the short term and then replacing them several years later with a newer more capable IP STB. Moving to a device that supports IP delivery also opens up many additional revenue opportunities for the operator such as subscribers upgrading to higher programming tiers and purchasing On Demand content. In addition, the migration to an all IP headend infrastructure allows for additional services to other devices in the home such as PCs, Tablets, Smart TVs, and Game Consoles eventually reducing the need for additional IP STBs.

There is an underlying technology convergence which serves to enable the migration to an IP DTA solution. These technologies include a combination of higher capacity DOCSIS CPEs and high powered IP STBs that support downloadable Conditional Access. In addition innovations in CMTS technology have driven DOCSIS downstreams costs down significantly over recent years.

An IP DTA solution is now a viable alternative to other methodologies to help solve the HFC spectrum exhaustion issue. The IP DTA solution also provides a leapfrog approach to IP services which is the ultimate destination for most operators.

References

1. "DOCSIS 3.0 Specifications", Cable Labs, http://www.cablelabs.com/cablemodem/specifications/specifications30.html
2. "Scaling the DOCSIS network for IPTV", Sangeeta Ramakrishnan, Cisco Systems, SCTE Emerging Technologies Conference 2009.
3. "Converged Cable Access Platform (CCAP)", Cable Labs Specification http://www.cablelabs.com/specifications/CM-TR-CCAP-V03-120511.pdf
4. "Variable Bitrate Video Services in DOCSIS 3.0 Network", Xiaomei Liu and Alon Bernstein, Cisco Systems, NCTA 2008

Application of 4K-QAM, LDPC and OFDM for Gbps Data Rates over HFC Plant

David John Urban
Comcast

Abstract

This paper describes the application of 4096-QAM, low density parity check codes (LDPC), and orthogonal frequency division multiplexing (OFDM) to transmission over a hybrid fiber coaxial cable plant (HFC). These techniques enable data rates of several Gbps.

A complete derivation of the equations used for log domain sum product LDPC decoding is provided. The reasons for selecting OFDM and 4K-QAM are described. Analysis of performance in the presence of noise taken from field measurements is made.

INTRODUCTION

A new physical layer is being developed for transmission of data over a HFC cable plant. The objective is to increase the HFC plant capacity.

A straightforward method to increase capacity over the HFC plant is to allocate more spectrum.The spectrum used for RF signals is typically 5-42 MHz in the upstream and 54-750 MHz in the downstream for HFC plants in the United States, although many variations of the upstream and downstream frequency allocations exist. The downstream signals can be extended to 1002 MHz to increase downstream capacity. The upstream can be increased to 5-85 MHz or even higher to increase upstream capacity. The capacity can be increased with higher spectral efficiency, the bits per second in a Hz of bandwidth. The capacity can be increased by packing more signals into existing spectrum, for example by reducing loss due to filter rolloffs, and adding more robust signals that can work in noisy parts of the spectrum.

Higher spectral density requires mapping more bits into a symbol at the expense of a higher signal to noise ratio (SNR) threshold. The new physical layer will map 12 bits per symbol on the downstream compared to the 8 bits per symbol used in DOCSIS 3.0. 4K-QAM (or 4096-QAM) is a modulation technique that has 4,096 points in the constellation diagram and each point is mapped to 12 bits. The new physical layer will map 10 bits per symbol using 1024-QAM (or 1K-QAM) on the upstream compared to 6 bits per symbol with 64-QAM used in DOCSIS 3.0.

INCREASING SPECTRAL EFFICIENCY

The three key methods for increasing the spectral efficiency of the HFC plant in the new physical layer are LDPC, OFDM and 4K-QAM. LDPC is a very efficient coding scheme. OFDM divides the spectrum into small subcarriers. 4K-QAM maps each symbol into 12 bits.

LDPC is a linear block code that utilizes a sparse parity check matrix. Since the parity check matrix is sparse, the code word can be very long with a reasonable calculation complexity. With a long code word size the parity check equations can have statistically significant samples. This allows LDPC codes to be demodulated with an iterative message passing algorithm based upon the log-likelihood ratio (LLR) for each bit of the code word.

The data rate is calculated by multiplying the spectral efficiency by the channel width, so the data rate can be doubled by doubling the channel width. A wide channel width results in a short symbol period; this increases the length of the adaptive equalizer needed to mitigate the inter-symbol interference due to reflections.

OFDM is a simpler method to increase plant capacity. OFDM divides the wide channel width into many subcarriers - a fast Fourier transform provides an efficient mathematical algorithm to divide the channel into many subcarriers. The subcarrier spectral width is narrow in the frequency domain. The narrow channel width subcarriers have a long time period; the OFDM symbol time is also very long. This has two benefits:

- Short in time impulse noise is averaged over a long symbol time.
- A guard interval can be added to efficiently eliminate inter-symbol interference due to echoes.

OFDM has another advantage: each subcarrier can be bit loaded differently to optimize for the changes in SNR over frequency.

The advantages of OFDM with an LDPC code allow for higher spectral efficiency. In addition, rather than mapping 6 or 8 bits to a symbol, 4K-QAM maps 12 bits to each data subcarrier in the OFDM symbol.

COMPARING DOCSIS TO THE NEW PHY

DOCSIS downstream signals conform to the ITU standard J.83B in North America. The symbol rate is 5.360537 MHz and the modulation is 256-QAM, 8 bits per symbol. It employs Reed-Solomon forward error correction with 122 information symbols and 128 code word symbols along with trellis coded modulation with an overall code rate of 19/20 is employed. Thus, the overall data rate is 38.8304 Mbps.

The occupied bandwidth and channel spacing is 6 MHz. Channel bonding 32 J.83B signals requires 192 MHz of spectrum and has a data rate of 1,242.6 Mbps. This paper will show that an OFDM signal using 4K-QAM and LDPC will have a data rate of 1.87 Gbps in 192 MHz occupied spectrum.

In DOCSIS, an adaptive equalizer with 24 taps covers 4.5 µs per 6 MHz channel; 32 bonded channels require 768 adaptive equalizer taps. OFDM is simpler: a single guard interval is added that can be ignored by the receiver. No knowledge of the channel impulse response is needed to eliminate inter-symbol interference with OFDM, so long as the guard interval is greater than the longest echo. It has been shown that in some cases only a small performance degradation results from shortening the guard interval to less than the longest echo. However, with so many sources of noise and interference that the cable operator has no control over, it is unwise to allow performance degradation that can be easily eliminated. The guard interval should be set to be longer than the longest expected echo.

TRANSMITTING AND RECEIVING SIGNALS

The process of creating a transmit signal in the new physical layer consists of several steps. The first step is to create a code word from information bits; this is done by multiplying the parity generator matrix with the information bits and then adding parity bits to the information bits to form a code word. Next, the bits of code words are mapped into symbols. With 4K-QAM 12 bits will be mapped into a real and imaginary discrete amplitude level taking on one of 64 different values. The symbols are assigned to subcarriers of different center frequencies, and an inverse fast Fourier transform (FFT) is used to create the time domain waveform of the useful part of the OFDM symbol. A cyclic prefix is added to the useful part of the OFDM symbol to provide a guard time.

After the signal goes through the channel, subjected to attenuation, noise and interference, the receiver separates the useful symbol period from the cyclic prefix. An FFT of the time domain useful symbol period outputs in-phase (I) and quadrature (Q) received vectors for each data subcarrier. From the received vectors, log-likelihood ratios (LLRs) of each bit in the code word are determined and passed on to the LDPC decoder. The LDPC decoder uses a message passing algorithm between check nodes and variable nodes to iteratively converge on the transmitted code word.

Received signals are degraded by three types of noise: thermal, ingress and impulse. Thermal noise is flat over frequency, determined by the receiver noise figure and the input level. Ingress noise is narrow in frequency and long in time, mostly due to radio interference. New cellular services based on LTE are particularly worrisome since they are typically 10 MHz wide and operate co-channel with HFC downstream frequencies. Impulse noise is narrow in time but wide in frequency; the most common source being electrical systems such as motors and lighting systems.

The downstream receive signals are dominated by thermal and ingress noise, while the upstream received signals have all three types. The downstream is higher in frequency, so cable attenuation is higher and thermal noise is more of a problem. Impulse noise can be observed at downstream frequencies, but tends to be less prevalent than ingress interference. Radio interference from broadcast, cellular, and public safety impact downstream signal quality.

On the upstream impulse noise has very high amplitudes below 20 MHz. Radio interference from shortwave, amateur, and citizen's band services are also major interferers.

NEW PHY DETAILS

The following sections of the paper describe the components of the new PHY: 4K-QAM, OFDM, and LDPC. The LDPC section includes a complete derivation of equations used for decoding an LDPC signal. After the LDPC section, a section on performance analysis provides results of simulations using field measured noise.

4K-QAM

Spectral efficiency is determined primarily by the number of bits that get mapped to a symbol. In 64-QAM 6 bits are mapped into a symbol while 256-QAM maps 8 bits. The higher spectral efficiency of 1024-QAM maps 10 bits to a symbol, and, taking it 2 bits further, 4K-QAM maps 12 bits into a symbol. A pretty steep penalty is incurred by increasing the number of bits per symbol in terms of increased SNR requirements.

In general, QPSK at 2 bits per symbol works well at a SNR of 12 dB. Each increase of 2 bits per symbol will incur a 6 dB SNR penalty. Thus, 16-QAM requires SNR of 18dB, 64-QAM requires SNR of 24dB, 256-QAM requires SNR of 30dB, 1024-QAM requires SNR of 36dB, and 4096-QAM requires SNR of 42dB. Now, things are not quite as dire as these numbers would lead you to believe. Advanced forward error correction techniques such as LDPC described in an upcoming section can lower the SNR threshold considerably. In a pure additive white Gaussian noise (AWGN) channel with error correction of (16200,14323), 4K-QAM can have a 10^{-8} bit error rate at 35.23 dB SNR.

4K-QAM has 4,096 constellation points with in-phase (I) and quadrature (Q) components amplitude modulated at 64 discrete values {-63,-61,…,-1,+1,…+63}. The constellation diagram for 4K-QAM is shown in Fig.1. A QAM signal can be represented in complex form as I+jQ where j is the square root of minus 1. I is the in-phase or real part of the complex baseband signal and Q is the quadrature or imaginary part of the complex baseband signal. The modulated signal is I(t)cos(2πft)+Q(t)sin(2πft).

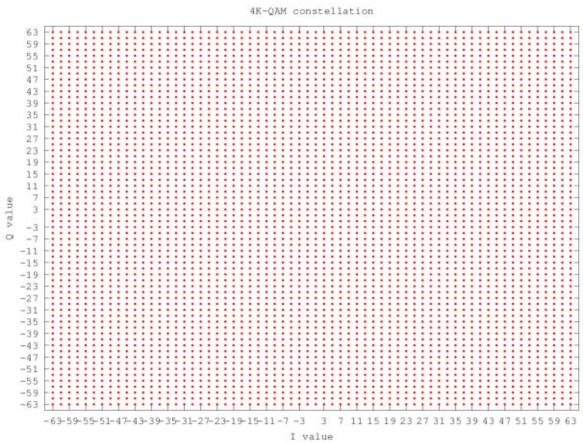

Fig. 1 4K-QAM Constellation Diagram.

The symbol error rate of 4K-QAM at an SNR of 42 dB is 1.29×10^{-3}. The symbol error rate can be calculated from the points in the constellation, M, and the linear SNR, E_s/N_0 using equation (1) [7]. A plot of the symbol error rate versus SNR for 4K-QAM is shown in Fig. 2. The symbol error rate is 22.9% for a signal to noise ratio of 35.23 dB at 4K-QAM.

$$P_{SEQAM} = 1 - \left\{1 - \left\{\left(1 - \frac{1}{\sqrt{M}}\right) \cdot erfc\left(\sqrt{\frac{3 \cdot E_s}{2 \cdot (M-1) \cdot N_0}}\right)\right\}\right\}^2 \quad (1)$$

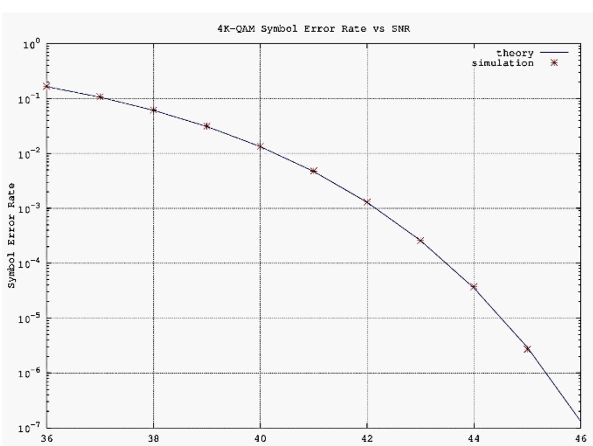

Fig. 2 Symbol Error Rate vs. SNR 4K-QAM

The constellation diagram of 4K-QAM with SNR of 43 dB is shown in Fig. 3. It can be seen that the points are well defined and that hard decision decoding of symbols will work well.

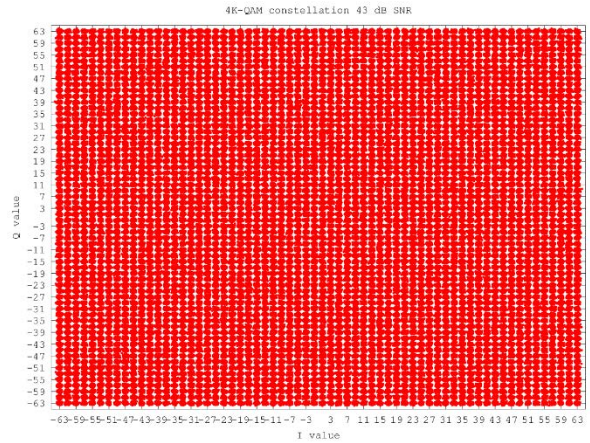

Fig. 3 4K-QAM with 43 dB SNR.

For a modulation of 4K-QAM the number of constellation points is 4096 and each symbol is mapped to 12 bits. 6 bits are mapped to an I value and the other 6 bits are mapped to a Q value. The I and Q levels are selected from the range of odd numbers between -63 and +63. The mapping of bits into I and Q levels for 4K-QAM is illustrated in Fig. 4 and listed in Table 1.

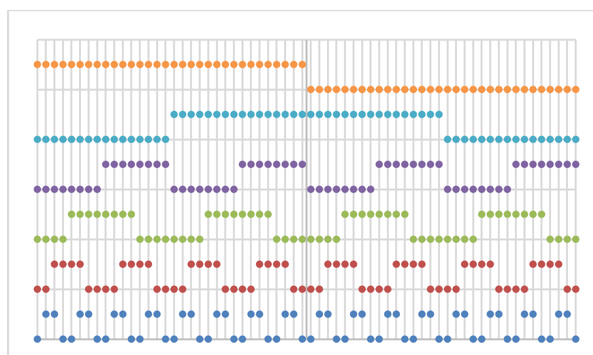

Fig. 4 Illustration of 4K-QAM bit mapping.

The least significant bit, shown as blue dots at the bottom of Fig.4, form a pattern 0,1,1,0 as the level goes from -63 to +63 in discrete odd steps along the vertical axis. The least significant bit is vulnerable to error across the entire level range.

On the other hand the most significant bit, represented in Fig. 4 by the orange dots at the top, are 1 for all negative amplitude levels and 0 for positive amplitude levels. The most significant bit is only in danger of error near the -1 to +1 level region. The bits are Gray coded [8], in that only one bit will change for each step in level.

With Gray coding [8] in general there will likely be only one or possibly two bit errors for each symbol error. A plot of bit error rate versus SNR for 4K-QAM is shown in Fig. 5. The bit error rate curve was determined by simulation using the bit mapping of Table 1, while the symbol error rate displays simulation results along with the closed form solution using equation (1). At 42 dB SNR the bit error rate for 4K-QAM is 10^{-4}. This is a low enough error rate that a Reed-Solomon forward error correction code with a very high coding rate can realize a 10^{-8} bit error rate. This is why we often think of 42 dB SNR to be a threshold level for 4K-QAM with RS coding.

A 42 dB SNR requirement is a bit too severe, especially when you consider that this number applies only to an AWGN channel with no implementation loss in the equipment and does not consider such factors as level variation over time and temperature, reflections in the cable plant, and interference from radio sources. A significant margin to the 42 dB threshold will be needed for real equipment in commercial service over a real cable plant. The margin needed will likely be at least 3 dB and perhaps more than 8 dB. That would mean that SNRs of 45 to 50 dB would be required for 4K-QAM.

I level	bit 6	bit 5	bit4	bit 3	bit 2	bit 1
-63	1	0	0	0	0	0
-61	1	0	0	0	0	1
-59	1	0	0	0	1	1
-57	1	0	0	0	1	0
-55	1	0	0	1	1	0
-53	1	0	0	1	1	1
-51	1	0	0	1	0	1
-49	1	0	0	1	0	0
-47	1	0	1	1	0	0
-45	1	0	1	1	0	1
-43	1	0	1	1	1	1
-41	1	0	1	1	1	0
-39	1	0	1	0	1	0
-37	1	0	1	0	1	1
-35	1	0	1	0	0	1
-33	1	0	1	0	0	0
-31	1	1	0	0	0	0
-29	1	1	0	0	0	1
-27	1	1	0	0	1	1
-25	1	1	0	0	1	0
-23	1	1	0	1	1	0
-21	1	1	0	1	1	1
-19	1	1	0	1	0	1
-17	1	1	0	1	0	0
-15	1	1	1	1	0	0
-13	1	1	1	1	0	1
-11	1	1	1	1	1	1
-9	1	1	1	1	1	0
-7	1	1	1	0	1	0
-5	1	1	1	0	1	1
-3	1	1	1	0	0	1
-1	1	1	1	0	0	0
1	0	1	0	0	0	0
3	0	1	0	0	0	1
5	0	1	0	0	1	1
7	0	1	0	0	1	0
9	0	1	0	1	1	0
11	0	1	0	1	1	1
13	0	1	0	1	0	1
15	0	1	0	1	0	0
17	0	1	1	1	0	0
19	0	1	1	1	0	1
21	0	1	1	1	1	1
23	0	1	1	1	1	0
25	0	1	1	0	1	0
27	0	1	1	0	1	1
29	0	1	1	0	0	1
31	0	1	1	0	0	0
33	0	0	0	0	0	0
35	0	0	0	0	0	1
37	0	0	0	0	1	1
39	0	0	0	0	1	0
41	0	0	0	1	1	0
43	0	0	0	1	1	1
45	0	0	0	1	0	1
47	0	0	0	1	0	0
49	0	0	1	1	0	0
51	0	0	1	1	0	1
53	0	0	1	1	1	1
55	0	0	1	1	1	0
57	0	0	1	0	1	0
59	0	0	1	0	1	1
61	0	0	1	0	0	1
63	0	0	1	0	0	0

Table 1. 4K-QAM Mapping Bits to I/Q Level.

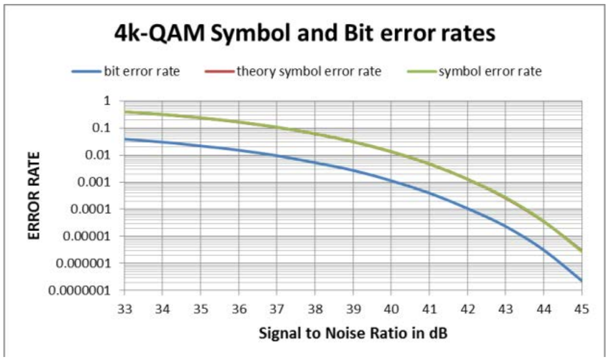

Fig. 5 Bit Error Rate for 4K-QAM.

Error correction techniques including concatenated outer coding, time and frequency, and bit interleaving will need to be employed to lower the SNR requirement in an AWGN channel to 35.23 dB. The code rate required to get to a 35.23 dB SNR threshold will be 88%; the 12 bits per symbol of 4K-QAM will then have a spectral efficiency of 10.5 bits/sec/Hz. This is a half bit higher than 1K-QAM which is 10 bits per symbol and has a general SNR threshold of 36 dB. This shows that while the proposed coding scheme is good, it is not so outlandishly good that it falls in the “too good to be true” category. We are expecting this coding scheme to get a half bit higher spectral efficiency at a 0.77 dB lower SNR than QAM modulation alone can do.

By far the most important element in our error correction scheme that will allows us to employ 4K-QAM modulation with an AWGN SNR threshold of 35.23 dB is LDPC. The key to LDPC decoding is an iterative message passing sum product method based on the log-likelihood of the bits in the code word given the received signal vector. This means that rather than the hard decoding used in determining the error rate in Fig 5, the output of the demodulator will be probabilities that a bit is a 1 or a 0. Fig. 3 shows that with a 43 dB SNR we can make a good estimate of the symbol that was sent and then determine the bits that were sent from the bit map table. In addition, Fig. 5 shows that this method provides a reasonable bit error rate. But now look at Fig. 6.

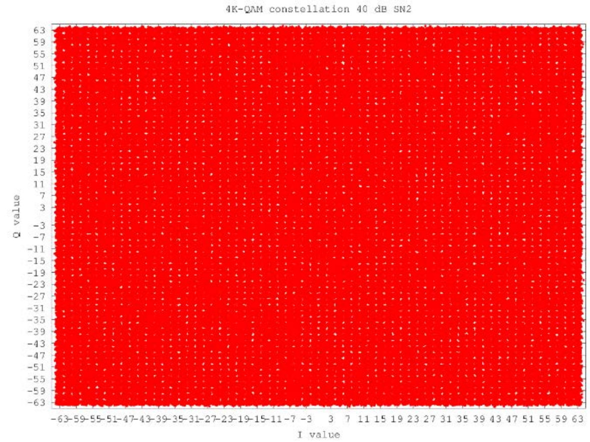

Fig. 6 4K-QAM Constellation at 40 dB SNR.

At 40 dB SNR the symbols in the constellation are becoming indistinguishable, as shown in Fig. 6. Showing a constellation of 4K-QAM at 35.23 dB SNR is not useful, since it would just be a solid red rectangular blob. The symbol error rate for 4K-QAM at an SNR of 35.23 dB is 21.9% and the bit error rate is 2%.

For every symbol that is received, an I and Q level is measured. From the measured I and Q levels the probability of a 1 having been transmitted or a zero can be calculated for each of the 12 bits mapped to a 4K-QAM symbol. Fig. 7 shows the R language code that was used to calculate the LLR of each bit in a 4K-QAM symbol as a function of the received amplitude of the quadrature components. The code runs through all possible receive levels. For each receive level the program calculates the distance to each of the 64 possible transmitted amplitude levels. The probability is calculated using the Gaussian probability density function. The probabilities of 1s are added and the probability of 0s are added for each of the 6 bits. The LLR of each bit is calculated by taking the natural logarithm of the ratio of the probability of a 0 divided by the probability of a 1.

```r
b=12; M=2.^b; m=1:sqrt(M); Iv=2*m-sqrt(M)-1;
map<-read.csv("map.csv"); Pavg=0;
for ( i1 in 1:sqrt(M) ) { for ( j1 in 1:sqrt(M) ) { Pavg=Pavg + Iv[i1]*Iv[i1]+Iv[j1]*Iv[j1] } }
Pavg=Pavg/M; SNR=35.23; SNRlin=10.^(SNR/10); N0=Pavg/SNRlin; sigma=sqrt(N0/2);
rx=seq(-65,65,length=4096);
p0=mat.or.vec(length(rx), ( ncol(map)-1)  ); p1=mat.or.vec(length(rx), ( ncol(map)-1)  );
for ( i1 in 1:length(rx) ) {
    for ( i2 in 1:nrow(map) ) {
        for ( j1 in 2:ncol(map) ) {
            if ( map[i2,j1]==0) {
                amplitude=1/(sqrt(2*pi)*sigma);
                numerator=-1*(map[i2,1]-rx[i1])^2;
                denominator=2*sigma^2;
                pd=amplitude*exp(numerator/denominator);
                p0[i1, (j1-1) ]=p0[i1,(j1-1)]+pd;
            }
            else {
                amplitude=1/(sqrt(2*pi)*sigma);
                numerator=-1*(map[i2,1]-rx[i1])^2;
                denominator=2*sigma^2;
                pd=amplitude*exp(numerator/denominator);
                p1[i1, (j1-1) ]=p1[i1,(j1-1)]+pd;
                }
                        }

                    }
 }
llr=log(p0)-log(p1);
plot(rx[],llr[,1],col="magenta",main="Log Likelihood Ratio 4k-QAM 35.23 dB SNR",xlab="received I value"
        ,ylab="LLR",type="l",cex.lab=2, cex.axis=2, cex.main=2, cex.sub=2, ylim=c(-100,100) );
lines(rx[],llr[,2],col="orange",type="l",cex.lab=2, cex.axis=2, cex.main=2, cex.sub=2);
lines(rx[],llr[,3],col="black",type="l",cex.lab=2, cex.axis=2, cex.main=2, cex.sub=2);
lines(rx[],llr[,4],col="red",type="l",cex.lab=2, cex.axis=2, cex.main=2, cex.sub=2);
lines(rx[],llr[,5],col="green",type="l",cex.lab=2, cex.axis=2, cex.main=2, cex.sub=2);
lines(rx[],llr[,6],col="blue",type="l",cex.lab=2, cex.axis=2, cex.main=2, cex.sub=2);
grid();
```
Fig. 7 R code for calculating 4K-QAM LLR.

The resulting plot of LLR for all 6 bits versus the I and Q of the receive level is shown in Fig. 8. The vertical axis has a unit of volts, although the scale of the voltage waveform will vary greatly at various stages of the transmission and reception process. The vertical axis is the natural logarithm of the ratio of the probability of a 0 divided by the probability of a 1 for each of the 6 bits mapped to a symbol. The vertical axis represents the received level scaled to match the discrete output levels of the symbol mapper {-63,-61,…,-1,+1,…,+63}. Note that blue curve is the least significant bit, it forms a uniform saw tooth pattern over the range of receive levels. As the bits become more significant the LLR peaks increase and the number of periodic repetitions decreases. The most significant bit shows up as a single line with a value of 0 for a received level of 0, off the charts positive value for receive levels above 10, and off the chart negative value for received levels below -10. Since a lower probability event contains more information than a high probability event, the vast majority of the information contained in the LLRs is found in the least significant bits.

In an effort to better understant the LLR, Fig. 9 zooms in on the least significant bit from receive level -65 to -45 and shows red dots to indicate a 0 or a 1. If the transmitted level is -63, -57, -55, -49, or -47, then the least significant bit is a 0. If the transmitted level is

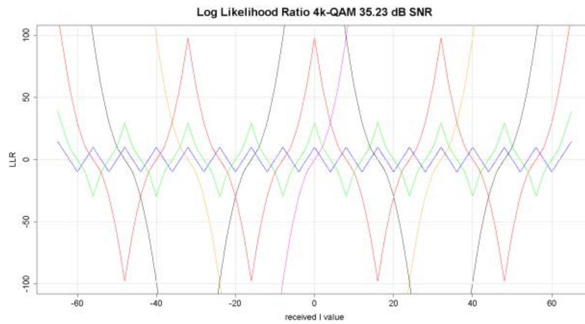

Fig. 8 4K-QAM LLR vs. receive level.

a -61, -59, -53, -51, -45 then the least significant bit is a 1. Consider the case where the received vector is -60. The transmit levels to the left and right both indicate a 1 was sent. At this point the LLR reaches a peak negative level of about -8. This indicates that there is a good likelihood that the transmitted bit was a 1. On the other hand, if the received level is -50, then the bit on the left is a 1 and the bit on the right is a 0. Both are equidistant so we really have no idea whether the transmitted bit was a 1 or a 0. In this case the LLR is 0. A positive LLR indicates a likelihood of a 0 and a negative LLR indicates the likelihood of a 1. The greater the magnitude of the LLR the greater the certainty.

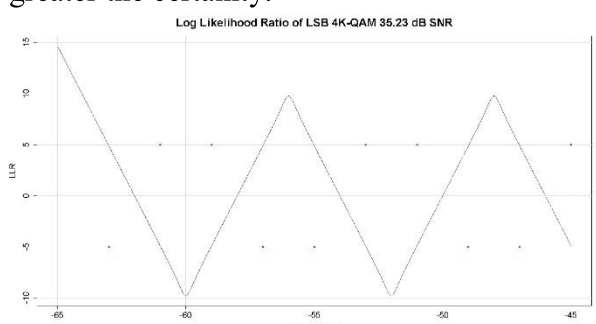

Fig. 9. 4K-QAM LLR of least significant bit.

As has been shown, the benefit of 4K-QAM is higher spectral efficiency - mapping a larger number of bits into a symbol. The penalty for high spectral efficiency is an increase in the SNR requirements. We have shown that we can map 12 bits into a symbol and use an error correction scheme with 88% coding rate to get 10.5 bits/sec/Hz with an SNR of 35.23 dB. The key to high spectral efficiency at a reasonable SNR is low density parity check codes with message passing iterative decoding based upon LLRs. We've shown how the LLRs that feed the LDPC decoder are found for 4K-QAM.

This still by itself does not give us our objective of high data rates. High spectral efficiency is only part of the equation. The data rate is the spectral efficiency times the channel width. To get a high data rate a very wide channel width is needed, even when the spectral efficiency is high. For this, OFDM is needed. A 2 Gbps data rate requires 200 MHz of channel width if the spectral efficiency is 10 bits/sec/Hz.

Description	Symbol	Value	Units
Guard Interval (Cyclic Prefix)	T_g	2.5	μs
Useful Symbol Time	T_u	40	μs
subcarrier frequency spacing	δf	25	kHz
FFT sampling frequency	$R_{sampling}$	204.8	MHz
FFT size	N_{FFT}	8,192	subcarriers
channel width	CW	190	MHz
available subcarriers	N_A	7,600	subcarriers
pilot spacing	δpilots	128	
number of pilots	pilots	60	subcarriers
data subcarriers	N_{data}	7,540	subcarriers
constellation points	M	4096	points
data subcarrier bit loading	b	12	bits
FEC information bits	K_{BCH}	14,232	bits
FEC codeword size	N_{LDPC}	16,200	bits
data rate	R_{data}	1,870.3	Mbps
signal to noise ratio threshold	SNR	35.23	dB

Table 2 Key Parameters of Downstream OFDM Signal.

OFDM

OFDM is a technique that greatly facilitates operating with very wide channel widths, while still remaining spectrally efficient, tolerating multi-path echoes without inter-symbol interference, dealing with narrow band interference, and averaging out the impact of short duration impulse noise by using a long symbol time.

The parameters of the downstream OFDM signal are listed in Table 2. The data rate is 1.8702 Gbps. 190 MHz of channel width is needed to get to such a high data rate. In addition, 1 MHz of frequency guard band is needed on both ends of the OFDM signal so that the total occupied bandwidth is 192 MHz, equivalent to 32 DOCSIS 3.0 downstream signals each occupying 6 MHz channel width. A DOCSIS 3.0 channel has a data rate of 38.8 Mbps for J.83 Annex B [12] so 32 bonded channels have a data rate of 1.24 Gbps. Thus, a data rate for OFDM of 1.87 Gbps is a capacity increase of 50.8% and an additional 630 Mbps of capacity when compared to 32 bonded DOCSIS channels in the same occupied spectrum.

OFDM helps in increasing the data rate in a number of ways. First, the channel width can be very large while the symbol period is also long. With a long symbol time, a guard interval can be added that is longer than the time delay of echoes in the channel without losing too much overall spectral efficiency. The individual subcarriers of the OFDM signal have a narrow frequency width, thus narrowband interference may only result in errors on a few subcarriers out of thousands. With a long OFDM symbol time, short time duration impulse noise can be averaged out over the entire symbol, lessening the impact. Selecting the OFDM parameters involves first determining the time delay of echoes in the channel due to reflections; this determines the guard interval. Then, a useful symbol time is chosen that is sufficiently large compared to the guard interval, so as to keep overall spectral efficiency high. The useful symbol time will determine the subcarrier frequency spacing so that an FFT size can be chosen that has a sampling frequency that is larger than the channel width. These steps are described in this section for the downstream signal with parameters shown in Table 2.

OFDM divides the channel width into thousands of small subcarriers. The inverse of the frequency spacing between subcarriers is equal to the useful symbol time. For a subcarrier spacing of 25 kHz the useful symbol time is 40 μs. A cyclic prefix or guard interval is added between symbols. With a cyclic prefix of 2.5 μs the symbol time is 42.5 μs.

The useful symbol time is the inverse of the subcarrier frequency spacing. A guard interval is added before the next symbol is sent; this guard interval is called a cyclic prefix, since it is constructed by adding a portion of the end of the time domain OFDM symbol waveform to the beginning. The guard interval is chosen to be longer than the time delay of echoes in order to prevent inter-symbol interference.

The longest echoes in the coaxial cable plant are due to express cables, which are used to reduce the number of cascaded amplifiers. Rather than feeding a far-off amplifier through a series of cascaded amplifiers, an express cable is connected from the fiber node directly to far-off amplifier. Express cables can be several thousand feet and they do not have taps which reduces the signal attenuation.

A 4.5 μs echo results from a signal traveling back and forth over a 1925 foot express cable. The attenuation for a 750 MHz signal traveling back and forth over a 0.875 inch diameter express cable of length 1925 feet is 23 dB as shown in Fig.10. There is less attenuation at lower frequencies; at 100 MHz the attenuation is only 8 dB. The amplitude of the 4.5 μs echo is determined by the attenuation of the express cable and the return loss of the devices at the end of the express cable. For a 10 dB return loss at the end devices, the amplitude of the echo will be -28 dBc at 100 MHz. Most fiber nodes and amplifiers have a minimum return loss specification of 16 dB. With a return loss of 16 dB for the devices at each end, the 4.5 μs echo will have an amplitude of -55 dBc at 750 MHz.

Analysis of cable modem equalizer taps and cable system maps found echoes that led to an estimate of the expected setting for the cyclic prefix (CP). A CP of 2.5 μs (trunk length=1070ft) should prove to be a universal setting. A CP of 5 μs (trunk length=2139ft)

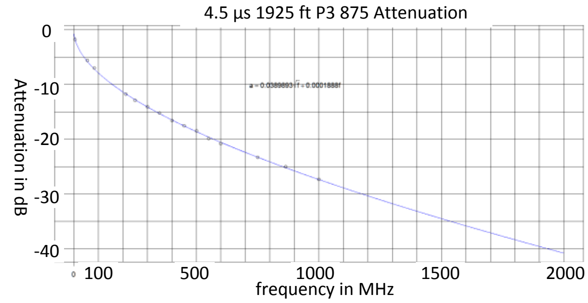

Fig. 10. Attenuation of Express Cables.

may on rare occasions be needed. It is possible that a 5 μs CP may prove more generally beneficial. A CP of 1 μs (trunk length=428ft) should work for most nodes and for most modems but may lead to some operational troubles.

Once the guard interval is determined that will prevent intersymbol interference in the presence of echoes, then a useful symbol time can be chosen. The useful symbol time should be long compared to the guard interval in order to avoid losing a large percentage of capacity. With a 2.5 μs guard interval, a useful symbol time of 40 μs is a reasonable choice.

When the useful symbol time is 40 μs, the subcarrier frequency spacing is 25 kHz. While a long useful symbol time relative to the cyclic prefix increases spectral efficiency, it also decreases the subcarrier frequency spacing and increases the size of the FFT. To avoid implementation loss and reduce complexity, the subcarrier frequency spacing should be made as wide and the FFT size should be made as small as the requirements allow.

Phase noise in oscillators and frequency synthesizers can degrade performance if the subcarrier frequency spacing is too small. The size of the FFT is the OFDM sampling frequency divided by the subcarrier frequency spacing. The OFDM sampling frequency must be larger than the channel width. In Table 2, the OFDM sampling frequency is 204.8 MHz, the channel width is 192 MHz, and the available subcarriers occupy 190 MHz of spectrum. If the subcarrier frequency spacing is too small relative to the channel width then the size of the FFT can get very large. With a sampling frequency of 204.8 MHz and a 25 kHz subcarrier frequency spacing, the FFT size is 8192. Limiting the FFT size to 8192 subcarriers keeps the complexity low while still providing a long enough useful symbol time.

One benefit of OFDM is that the conversion from modulated subcarriers to a time domain waveform can be computed using an FFT, algorithm. Likewise, the receiver can use an FFT algorithm to convert the time domain signal into the modulated frequency domain subcarriers. The FFT algorithm is computationally efficient. The number of complex computations of an FFT is on the order of Nlog(N) rather than the N^2 computations of a discrete Fourier transform. The FFT algorithm requires that the size of the FFT is 2^ν, where ν is an integer. The downstream FFT size is 8192, so ν=13. The occupied channel width is 192 MHz with a 1 MHz frequency guard band on each side, so the available subcarriers must fit in 190 MHz.

4K-QAM requires and SNR of 35.23 dB. The channel for this SNR threshold is AWGN, additive white Gaussian noise, which means that the noise is flat over the full channel width and the signal is steady. The forward error correction used is a concatenated outer code along with a (16200,14400) LDPC inner code. 14,232 information bits go into the outer encoder to form a 16,200 bit LDPC codeword.

Thermal noise is always present in any electronic device and it is flat over frequency. The electrons in a resistor moving about randomly due to the energy at temperatures above absolute zero create thermal noise. At room temperature, the thermal noise is -174 dBm/Hz. Adding 30+10*log10(75) will convert from dBm to dBmV. A noise bandwidth of 192 MHz will add 82.8 dB. The noise figure of the receiver will add to the noise level. With a 10 dB receiver noise figure and a 192 MHz noise bandwidth, the noise level is -34.2 dBmV. A 40 dB SNR will thus require an input level of +7.5 dBmV if the receiver noise figure is 10 dB. This input level within a 6 MHz channel width is -7.5 dBmV.

The cable attenuation at a frequency of 1000 MHz for 150 feet of F6 drop cable is 10 dB, as shown in Fig. 11. The attenuation for 500 feet of P3 500 trunk cable at 1000 MHz is 12 dB. With a 4-way splitter having 7 dB loss, an additional 100 feet of in home cable having a loss of 7 dB at 1000 MHz, five in line directional couplers each having 2 dB of through loss, and an 11 dB last tap, the required output level from the amplifier is +49.5 dBmV per 6 MHz.

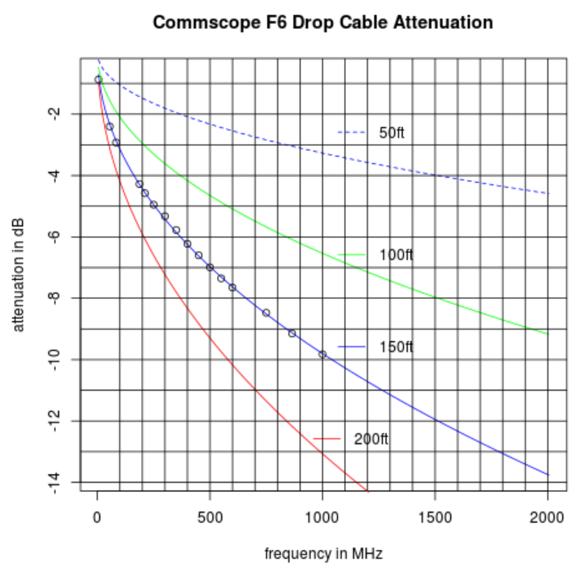

Fig. 11 Drop cable attenuation.

An analysis of input levels to cable modems found that the average level is around 0 dBmV with a standard deviation of around 6 dB. Thus, about 85% of cable modems have an input greater than -6 dBmV per 6 MHz, enough for 4K-QAM. The link budget of the cable plant is sufficient to provide most modems with an input level high enough to support 4K – QAM.

To form an OFDM signal the output bits from the code word generator are subdivided

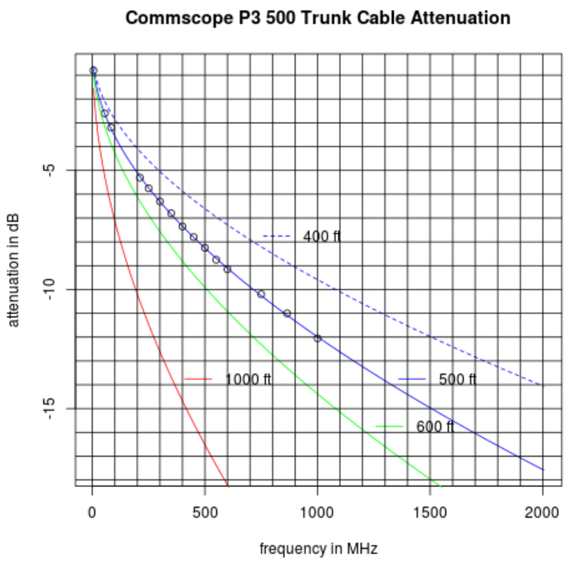

Fig. 12 Trunk cable attenuation.

to match the number of bits into each subcarrier. Downstream signals operating in the flat part of the spectrum in the absence of radio frequency interference will map each data subcarrier into same number of bits, typically 12 bits for each data subcarrier. The 16,200 bits from the code word generator are mapped into 1,350 data subcarriers for 4K-QAM. The next section describes the coding of the OFDM signal, LDPC.

LDPC

LDPC, is a linear block code. The codeword consists of information bits and parity bits, created by multiplying the information bits by a codeword generation matrix. Systematic codewords consist of the information bits followed by parity bits. As shown in equation (2), the codeword generation matrix of a systematic code, G, consists of an identity matrix and a parity generation matrix, denoted as P, with the codeword denoted by c and the information bits denoted by x.

$$c = x \cdot G = x \cdot [\, I_k \,|P] = [\, x\; p\,] \qquad (2)$$

The codeword, c, is a matrix with 1 row and nc columns, where nc is the number of bits in a codeword. The downstream LDPC codeword has a length 16,200 bits. The upstream short packet code word has a length of 1,120 bits. The upstream medium packet code word has a length of 5,940 bits. The upstream long packet code word has a length of 16,200 bits.

The information bits can be represented as a column vector having one row and kc columns denoted as x in equation (2). The codeword generation matrix, G, has kc rows and nc columns. The parity bits are represented in equation (2) as a vector, p, with one row and mc columns generated by the matrix P, which has kc rows and mc columns. The downstream LDPC codeword has 14,400 information bits, denoted by the variable kc. The downstream LDPC codeword has 1,800 parity bits, denoted by the variable mc=nc-kc. The upstream LDPC short codeword with a length of 1,120 bits has 840 information bits and 280 parity bits. The upstream medium size codeword with a length of 5,940 bits has 5,040 information bits and 900 parity bits. The upstream long codeword has the same parameters as the downstream code word, nc=16200, kc=14400, mc=1800.

$$\begin{gathered}
[H] \cdot [G^T] = [0] \\
[(mc\ by\ nc)] \cdot [(nc\ by\ kc)] = [(mc\ by\ kc)] \\
[A\,|B] \cdot \begin{bmatrix} I_k \\ P^T \end{bmatrix} = [0] \\
A \cdot I_k + B \cdot P^T = 0 \\
B \cdot P^T = -A \\
P^T = B^{-1} \cdot (-A) \\
P = (B^{-1} \cdot (-A))^T
\end{gathered} \qquad (3)$$

Equation (3) can be used to calculate the parity bit generator matrix, P, from a decomposition of the parity check matrix, H. The parity check matrix can be described with a base matrix and a lifting factor. The base matrix can be expanded by a series of circular shifts of identity matrices to form the parity check matrix. It is then necessary to determine the parity generator matrix from the parity check matrix. The equation is shown in (3) and Fig. 13 shows the Octave code used in simulations to form the parity check matrix from a base matrix and then determine the parity generator matrix for the upstream long code. The base matrix, Hbm, is a 5 row and 45 column matrix whose elements indicate the amount of right circular shift to apply to a 360 by 360 identity matrix. If an element of the base matrix is -1 then a 360 by 360 zero matrix is the sub-matrix in this spot. 5 rows and 45 columns of these 360 by 360 sub-matrices form the parity check matrix which has 16,200 columns and 1,800 rows. This method of formulating the parity check matrix is called quasi-cyclic, denoted QC-LDPC.

A diagram showing the sparse matrix structure of the parity check matrix for the upstream long code word is shown in Fig. 14. This illustrates that some sub-matrices are all zero 360 by 360 matrices, while others are circularly shifted 360 by 360 identity matrices. The secret sauce of LDPC codes is that long codes can be used in order to gain statistically significant channel information while keeping the computational complexity within reason since the parity check matrix is sparse.

HPgen16200.m

```
tic()
Hbm=load("Hbm16200.txt");
z=360;mb=5;nb=45;
nc=16200;mc=1800;kc=14400;
H=[];
for i1=1:rows(Hbm)
        Hrow=[];
        for j1=1:columns(Hbm)
                if ( Hbm(i1,j1) == -1 ) Z=zeros(z);
                else Z=shift(eye(z),-1*Hbm(i1,j1) );
                endif
                Hrow=[Hrow Z];
                Z=[];
        endfor
        H=[H;Hrow];
endfor
H=gf(H);
A=H(:,1:kc);
B=H(:,(kc+1):nc);
Binv=inv(B);
PT=Binv*A;
P=PT';
Pd=double(P.x);
Hs=sparse(double(H.x));
save Pdup16200.mat Pd;
save Hsup16200.mat Hs;
toc()
```

Fig. 13 Octave Code for Generating the Parity Check Matrix and Parity Bit Generator Matrix.

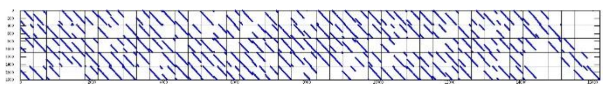

Fig. 14 Spy Diagram of 16200 bit upstream parity check matrix.

Before diving into the equations used in LDPC decoding, let's review how well the chosen codes work. The downstream (16200,14400) LDPC code can operate at 10^{-8} bit error rate using 4K-QAM in an AWGN channel with an SNR of 35.23 dB. The upstream short LDPC code has a bit error rate of 10^{-8} at an SNR of 28.77 dB for 1K-QAM. The upstream medium length LDPC code has a bit error rate of 10^{-8} at an SNR of 29.10 dB for 1k-QAM. The upstream long LDPC code has a bit error ratio of 10^{-8} at an SNR of 29.71 dB for 1k-QAM.

The spectrum analyzer measurement of a 5-42 MHz upstream OFDM signal taken at the 1K-QAM SNR threshold of 29.71 dB is shown in Fig. 15. The constellation diagram for 1K-QAM with an SNR of 29.71 dB is shown in Fig.16.

The constellation diagram at the signal to noise ratio threshold clearly reveals that hard decision decoding, that is making a decision about which symbol was sent based upon the symbol boundaries, will not work. Instead, the receiver will utilize the bit mapping table and the received signal vector to determine the LLR for each bit of the codeword. For a 16,200 bit codeword the decoder will be passed 16,200 LLRs, which will be the probability of the bit being a 1 or a 0 based upon the measured receive vector. If the LLR is greater than 0 then the initial best estimate for the bit is a 0, otherwise the best initial guess is a 1. The parity check matrix can be multiplied by the transpose of the estimate of the codeword to form the syndrome. If the syndrome is a zero vector, then the estimate is a code word.

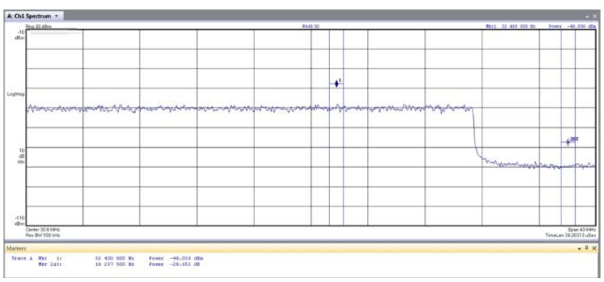

Fig. 15 Upstream 5-42 MHz Spectrum Plot.

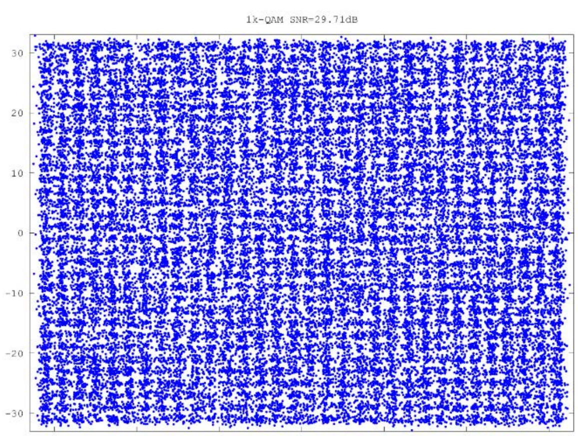

Fig. 16. The Constellation Diagram for 1K-QAM with an SNR of 29.71 dB.

Soft Decision Decoding

The equations used in soft decision decoding for LDPC are discussed in this section. The length of the codeword is denoted nc. A codeword length of 16,200 bits is used for both upstream and downstream signals with additional codeword lengths available for upstream signals. The soft decision decoder is fed the probability of each codeword bit being a 1 or a 0 given the received signal vector. The decoder has a number of parity check equations, here denoted by the variable mc. The number of parity check equations for downstream signals is 1,800. The columns of the parity check matrix can be thought of as variable nodes; the rows of the parity check matrix can be thought of as check nodes.

The soft decision decoder passes messages back and forth between variable nodes and check nodes, iteratively converging on the correct codeword. The parity check matrix is a sparse matrix so that most of the elements of the matrix are 0. The elements in the parity check matrix that are 1 represent a connection between a variable node and a check node. If element $h_{i,j}$ of the parity check matrix H is equal to 1 ($h_{i,j}=1$), then variable node c_i is connected check node f_j. For the downstream codeword, i will range from 1,2,…1800 and j will range from 1,2,…,16200. The connections between variable nodes and check nodes can be represented by a Tanner graph [2]. An example of the connections for a check node is shown in Fig. 17 and for a variable node in Fig.18.

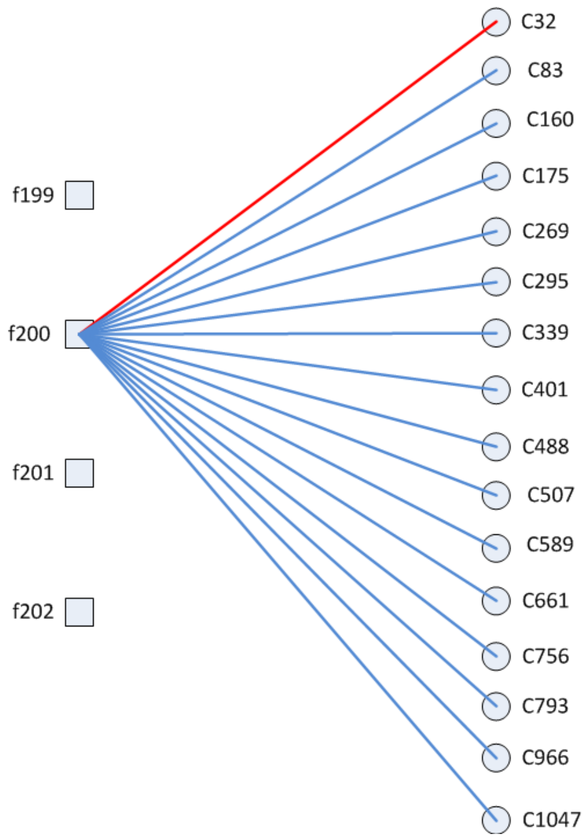

Fig. 17 Tanner Graph for Check Node 200 for LDPC Code (1120,840).

Fig. 17 illustrates the response messages that check node f200 will send to the 16 variable nodes connected to it. This graphical representation is called a Tanner graph [2]. The response message that check node f200 sends to variable node c32 is represented by the red line in Fig. 17. The message that is sent will be the probability that there are an even number of ones on the variable nodes connected with blue lines in Fig. 17. This probability is determined from the query messages sent from these variable nodes in the previous iteration.

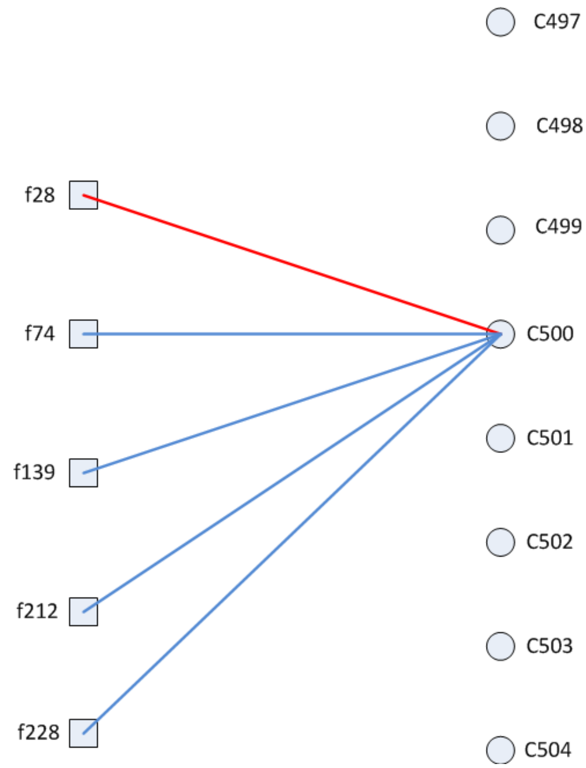

Fig.18 Tanner Graph of the Connections to Variable Node C500 for the LDPC Code (1140,840).

Fig. 18 shows the Tanner graph of the connections to variable node C500. The red line illustrates the response message from variable node C500 to check node f28. The response message will utilize only extrinsic information. The query message from f28 to C500 in the previous iteration will not be used in formulating the response message, since this would constitute passing back intrinsic information. The response message will be the product of five probabilities: the probability of a one given for variable node c500 given the channel received vector, and the query messages from check node f74, f139, f212, and f228 to variable node c500 from the last iteration.

The process of LDPC decoding begins with the LLRs from the demodulator. A vector is received from the channel and the demodulator must determine for each bit of the codeword the probability of a 1 or 0. The nc variable nodes send the LLRs to their check nodes. As shown in Fig. 18, each variable node is connected to a small number of check nodes. The check nodes will receive a query message from each of the variable nodes connected to it. What is the first message that a variable node should send to a check node? The first query message that variable node c_i will send to check node f_j will be the probability that the codeword bit is a 1 and the probability that the code word bit is a 0 based upon the received signal.

$$q_{ij}^{0}(1) = \Pr[c_i = 1|\vec{y}] = P_i \quad (4)$$
$$q_{ij}^{0}(0) = \Pr[c_i = 0|\vec{y}] = 1 - P_i \quad (5)$$
$$i = 1,2,\dots,nc$$
$$j = 1,2,\dots,mc$$

$$q_{ij}^{0} = \ln\left(\frac{q_{ij}^{0}(0)}{q_{ij}^{0}(1)}\right) = \ln\left(\frac{1-P_i}{P_i}\right) = l_i \quad (6)$$

This is shown in equation (4), (5), (6). The probability that the bit sent at variable node c_i is a 1, given the channel received vector y, is denoted as P_i. The query message sent from variable node c_i to check node f_j at iteration 0 is denoted q_{ij}^{0}. Since the probability of two events occurring is the multiplication of the probability of each event [5]; while working in the probability domain this calculation involves a lot of multiplication. The multiplications can be turned into additions by working in the log domain. Equation (6) shows the query message from variable node c_i to check node f_j at iteration 0 when working in the log domain, denoted as q_{ij}^{0}. This is the channel LLR for each variable node.

The check nodes receive messages from the variable nodes and must form a reply message. The parity check equation (7) tells us that there are an even number of 1s on variable nodes that are connected to a check node. The bit values of the variable nodes connected to a check node must add in GF(2) to zero, and they must add modulo 2 to zero. The exclusive OR function of all the bits on variable nodes connected to a check node must equal zero. All four of these ways of thinking (addition in a Galois field of 2, exclusive OR, modulo 2 arithmetic, even number of ones) are equivalent. It just depends on your level of comfort with finite linear algebra, Boolean equations, modulo arithmetic, logical functions.

$$H \cdot c^T = 0 \quad (7)$$

The messages sent between check nodes and variable nodes and vice versa should consist of extrinsic, and not intrinsic, information. We do not want the same messages passed back and forth between nodes without getting any closer to a correct code word. Thus, the reply message that check node f_j sends to variable node c_i at iteration l which is denoted as $r_{ji}^{l}(0)$ is the probability that there are an even number of 1s on all of the variable nodes connected to f_j except for the last query message sent from c_j.

Jumping ahead a bit, the response message from check nodes to variable nodes is given in equation (12). If you're curious as to how this equation is arrived at consider a couple simple cases. First, consider the probability of an even number of 1s from two variable nodes, given that q_1 is the probability of a one on variable node $c_1$as shown in equation (8).

$$\begin{aligned}\Pr[c_1 \oplus c_2] &= q_1 \cdot q_2 + (1-q_1)\cdot(1-q_2)\\ &= 1 - q_1 - q_2 + 2\cdot q_1\cdot q_2\\ &= \frac{1}{2} + \frac{1}{2}\cdot(1 - 2\cdot q_1 - 2\cdot q_2 + 4\cdot q_1\cdot q_2)\\ &= \frac{1}{2} + \frac{1}{2}\cdot(1-2\cdot q_1)\cdot(1-2\cdot q_2) = q\end{aligned} \quad (8)$$

Likewise, the probability of an even number of 1s from three nodes can be derived. The probability of an even number of 1s from three nodes is the probability of an odd number of 1s from two nodes, which has been found to be (1-q), and the probability of a 1 for the third node. This is shown in equation (9) and (10) and generalized in equation (11).

$$\Pr[c_1 \oplus c_2 \oplus c_3] = \frac{1}{2} + \frac{1}{2} \cdot \big(1 - 2 \cdot (1 - q)\big) \cdot (1 - 2 \cdot q_3) = \tfrac{1}{2} + \tfrac{1}{2} \cdot (2 \cdot q - 1) \cdot (1 - 2 \cdot q_3) \tag{9}$$

From equation (8) we have.

$$2 \cdot q - 1 = (1 - 2 \cdot q_1) \cdot (1 - 2 \cdot q_2)$$

$$\Pr[c_1 \oplus c_2 \oplus c_3] = \frac{1}{2} + \frac{1}{2} \cdot (1 - 2 \cdot q_1) \cdot (1 - 2 \cdot q_2) \cdot (1 - 2 \cdot q_3) \tag{10}$$

In general

$$\Pr[c_1 \oplus c_2 \oplus \cdots \oplus c_n] = \frac{1}{2} + \frac{1}{2} \cdot \prod_{i=1}^{n} (1 - 2 \cdot q_i) \tag{11}$$

A query message is sent from each variable node to all of the check nodes connected to it. This query message is denoted at $q_{ij}^{(l-1)}$, the message sent from variable node c_i to check node f_j at iteration (l-1). These query messages will be used to formulate the response messages.

The equation for the probability for an even number of 1s for n nodes given the probability of a 1 for each node is applied to determine the message that each check node sends to a connected variable node. Equation (12) is the formula used to calculate the response message $r_{ji}^{l}(0)$ from check node f_j to variable node c_i at iteration l, indicating the probability of a zero at this variable node using the query message from all the other variable nodes from the previous iteration.

$$r_{ji}^{l}(0) = \frac{1}{2} + \frac{1}{2} \cdot \prod_{i' \in V_{j \backslash i}} \left(1 - 2 \cdot q_{i'j}^{l-1}(1)\right) \tag{12}$$

The set of variable nodes that are connected to check node f_j is V_j. The set of all variable nodes that are connected to check node f_j, with the exception of variable node c_i, is denoted by $V_{j\backslash i}$.

The multiplications required to determine the response messages sent from the check nodes to the variable nodes can be replaced with additions by working in the log domain. The natural logarithm is plotted in Fig. 19, which reminds us of some fundamental properties: ln(1)=0, the argument of the natural logarithm must be positive and not zero; the natural logarithm is negative when the argument is less than one and positive when the argument is greater than one. It is helpful to visualize the curve of the natural logarithm when trouble shooting LDPC decoding algorithms, since violating these constraints will result in debugging errors.

Equation (13) defines the log domain query message from check node f_j to variable node c_i at the previous iteration, l-1. Through the algebraic manipulations shown in equations (14), (15),(16),(17),(18), equation (19) is arrived at, which allows the use of the hyperbolic tangent function.

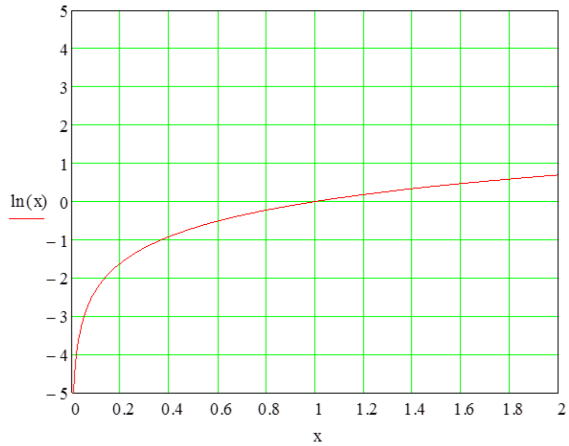

Fig. 19 Plot of the Natural Logarithm.

Fig. 20 shows a plot of the hyperbolic tangent function in order to remind us of some fundamental properties of this function. The tanh(x) function can have either positive or negative arguments, including zero. The tanh(x) function is an odd function. The output of tanh(x) is between +1 and -1. Equation (20) defines the hyperbolic tangent function, and equation (21) manipulates the hyperbolic tangent function into a form that can easily be applied to the equation for response messages.

$$q_{ij}^{l-1} = \ln\left(\frac{q_{ij}^{l-1}(0)}{q_{ij}^{l-1}(1)}\right) \tag{13}$$

$$e^{q_{ij}^{l-1}} = \frac{q_{ij}^{l-1}(0)}{q_{ij}^{l-1}(1)} = \frac{1 - q_{ij}^{l-1}(1)}{q_{ij}^{l-1}(1)} \tag{14}$$

$$q_{ij}^{l-1}(1) \cdot e^{q_{ij}^{l-1}} = 1 - q_{ij}^{l-1}(1) \tag{15}$$

$$q_{ij}^{l-1}(1) \cdot e^{q_{ij}^{l-1}} + q_{ij}^{l-1}(1) = 1 \tag{16}$$

$$q_{ij}^{l-1}(1) = \frac{1}{e^{q_{ij}^{l-1}} + 1} \tag{17}$$

$$1 - 2 \cdot q_{ij}^{l-1}(1) = \frac{e^{q_{ij}^{l-1}} + 1}{e^{q_{ij}^{l-1}} + 1} - 2 \cdot \frac{1}{e^{q_{ij}^{l-1}} + 1} \tag{18}$$

$$1 - 2 \cdot q_{ij}^{l-1}(1) = \frac{e^{q_{ij}^{l-1}} - 1}{e^{q_{ij}^{l-1}} + 1} \tag{19}$$

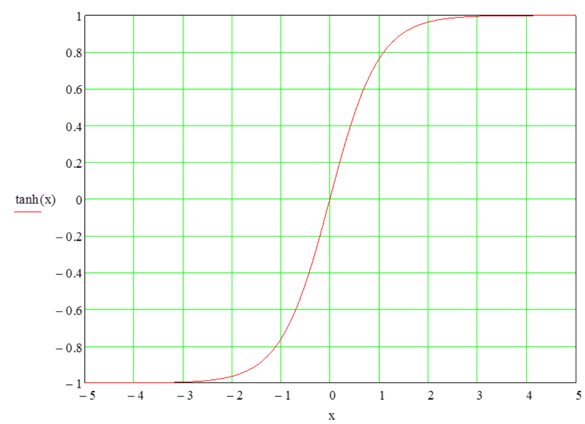

Fig. 20. Plot of the Hyperbolic Tangent.

$$\tanh(x) = \frac{\sinh(x)}{\cosh(x)} = \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}} \tag{20}$$

$$\tanh\left(\frac{x}{2}\right) = \frac{e^{\frac{x}{2}}}{e^{\frac{x}{2}}} \cdot \frac{e^{\frac{x}{2}} - e^{-\frac{x}{2}}}{e^{\frac{x}{2}} + e^{\frac{x}{2}}} = \frac{e^{x} - 1}{e^{x} + 1} \tag{21}$$

Substituting equation (21) into equation (19) gives equation (22).

$$\therefore\ 1 - 2 \cdot q_{ij}^{l-1}(1) = \tanh\left(\frac{q_{ij}^{l-1}}{2}\right) \tag{22}$$

Equation (22) allows us to use the hyperbolic tangent function in our equation for the response message (12) which is shown in equation (23).

$$r_{ji}^{l}(0) = \frac{1}{2} + \frac{1}{2} \cdot \prod_{i' \in V_{j \backslash i}} \tanh\left(\frac{q_{i'j}^{l-1}}{2}\right) \tag{23}$$

Now that we have an equation for the response message from a check node to a variable node that utilizes the hyperbolic tangent of the query node messages from the previous iteration in the log domain, it's time to convert the response messages to the log domain as shown in equation (24). Equation (24) defines the log domain representation of the response message from check node f_j to variable node c_i at iteration l.

$$r_{ji}^{l} = \ln\left(\frac{r_{ji}^{l}(0)}{r_{ji}^{l}(1)}\right) = \ln\left(\frac{r_{ji}^{l}(0)}{1 - r_{ji}^{l}(0)}\right) \tag{24}$$

Substituting equation (23) into equation (24) gives equation (25) which calculates the log domain response messages from check node f_j to variable node c_i at iteration l given the previous iteration log domain query messages from the variable nodes.

$$r_{ji}^{l} = \ln\left(\frac{\frac{1}{2} + \frac{1}{2} \cdot \prod_{i' \in V_{j \backslash i}} \tanh\left(\frac{q_{i'j}^{l-1}}{2}\right)}{\frac{1}{2} - \frac{1}{2} \cdot \prod_{i' \in V_{j \backslash i}} \tanh\left(\frac{q_{i'j}^{l-1}}{2}\right)}\right) \tag{25}$$

The numerator and denominator can be reversed by using the property of the natural logarithm shown in equation (26) resulting in equation (27).

$$\begin{aligned} \ln\left(\frac{a}{b}\right) &= \ln(a) - \ln(b) \\ &= -1 \cdot (\ln(b) - \ln(a)) \\ &= -\ln\left(\frac{b}{a}\right) \end{aligned} \tag{26}$$

$$\therefore r_{ji}^{l} = -\ln\left(\frac{1 - \prod_{i' \in V_{j \backslash i}} tanh\left(\frac{q_{i'j}^{l-1}}{2}\right)}{1 + \prod_{i' \in V_{j \backslash i}} tanh\left(\frac{q_{i'j}^{l-1}}{2}\right)}\right) \tag{27}$$

Equation (27) is in a form that can be matched with the hyperbolic tangent function shown in (28).

$$\tanh\left(\frac{x}{2}\right) = \frac{e^{\frac{-x}{2}}}{e^{\frac{-x}{2}}} \cdot \frac{e^{\frac{x}{2}} - e^{-\frac{x}{2}}}{e^{\frac{x}{2}} + e^{\frac{x}{2}}} = \frac{1 - e^{-x}}{1 + e^{-x}} \tag{28}$$

The response message equation can be simplified by breaking up the query messages into a sign part denoted by the alpha variable and a magnitude part denoted by beta, as defined in equation (29).

$$let\ \alpha_{ij}^{l-1} = sign\{q_{ij}^{l-1}\}\ and\ \beta_{ij}^{l-1} = \left|q_{ij}^{l-1}\right| \tag{29}$$

The property of the hyperbolic tangent function needed to simplify parts of equation (27) is shown in equation (30).

$$\begin{aligned} \tanh(-x) &= \frac{e^{-x} - e^{+x}}{e^{-x} + e^{+x}} = -1 \cdot \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}} \\ &= -\tanh(x) \end{aligned} \tag{30}$$

Equation (29) and (30) yields equation (31) which can be used to break apart the products in the numerator and denominator of equation (27) as shown in equation (32).

$$\tanh\left(\frac{q_{i'j}^{l-1}}{2}\right) = \alpha_{i'j}^{l-1} \cdot \tanh\left(\beta_{i'j}^{l-1}\right) \tag{31}$$

$$\prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{q_{i'j}^{l-1}}{2}\right) = \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} \cdot \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right) \tag{32}$$

There are two possible cases to consider in applying equation (32) into equation (27); one where the product of the signs of the query messages is positive and another where it is negative. These cases are shown in equation (33) and (34).

$$if \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} = 1 \; then \; r_{ji}^{l} = -\ln\left(\frac{1 - \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}{1 + \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}\right) \tag{33}$$

$$if \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} = -1 \; then \; r_{ji}^{l} = -\ln\left(\frac{1 + \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}{1 - \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}\right) = \ln\left(\frac{1 - \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}{1 + \prod_{i' \in V_{j\backslash i}} \tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}\right) \tag{34}$$

In both of these cases the response message can be expressed as shown in equation (35).

$$r_{ji}^{l} = \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} \cdot (-1) \cdot \ln\left(\frac{1 - \prod_{i' \in V_{j\backslash i}} tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}{1 + \prod_{i' \in V_{j\backslash i}} tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)}\right) \tag{35}$$

The substitution defined in equation (36) applied to equation (35) yields equation (37).

$$let \; e^{-x} = \prod_{i' \in V_{j\backslash i}} tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right) \tag{36}$$

$$then \; r_{ji}^{l} = \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} \cdot (-1) \cdot \ln\left(\frac{1 - e^{-x}}{1 + e^{-x}}\right) \tag{37}$$

Applying the definition of the hyperbolic tangent function shown in equation (28) to equation (37) allows us to further simplify the response equation as shown in equation (38).

$$r_{ji}^{l} = \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} \cdot (-1) \cdot \ln\left(\tanh\left(\frac{x}{2}\right)\right) \tag{38}$$

Returning to the definition of x found in equation (36), x can be expressed as shown in equation (39) and (40). Equation (40) can then be substituted into equation (38) to yield (41).

$$x = -1 \cdot \ln\left(\prod_{i' \in V_{j\backslash i}} tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)\right) \tag{39}$$

$$x = -1 \cdot \sum_{i' \in V_{j\backslash i}} \ln\left(tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)\right)$$

(40)

$$r_{ji}^{l} = \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} \cdot (-1) \cdot \ln\left(\tanh\left(\frac{-1 \cdot \sum_{i' \in V_{j\backslash i}} \ln\left(tanh\left(\frac{\beta_{i'j}^{l-1}}{2}\right)\right)}{2}\right)\right)$$

(41)

A new function can be defined shown in equation (42) which simplifies the response message equation as shown in equation (43). Equation (43) was used in log domain decoding simulations in the paper. A plot of the newly defined function in equation (42) is shown in Fig. 21.

$$let\ \varphi(x) = -\ln\left(tanh\left(\frac{x}{2}\right)\right)$$

(42)

$$r_{ji}^{l} = \prod_{i' \in V_{j\backslash i}} \alpha_{i'j}^{l-1} \cdot \varphi\left(\sum_{i' \in V_{j\backslash i}} \varphi\left(\beta_{i'j}^{l-1}\right)\right)$$

(43)

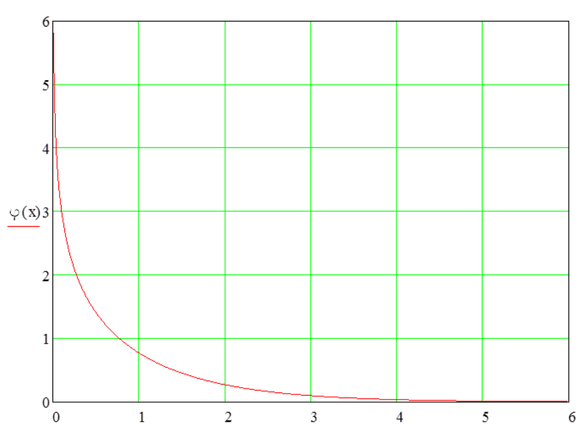

Fig. 21 Plot of phi(x)=-ln(tanh(x/2)).

Once response messages have been sent from the check nodes to the variable nodes, a new round of query messages will be sent from the variable nodes to the check nodes. What message should the variable node return to a connected check node? The first criteria is to pass extrinsic information, so variable node c_i does not to pass $r_{ij}^{(l-1)}$ back to check node f_j. This is the response message that just came from the check node so no progress is made in passing it right back. This response message should excluded.

The query message q_{ij}^{l} from variable node c_i to check node f_j at iteration l should use the response messages $r_{j'i}^{(l-1)}$, where j' does not equal j, as well as the probability for variable node c_i, determined from the channel received vector P_i. Since each of these messages is an estimate of probability, the probability of many events is determined by multiplying the probability of each event [5]. Then, we make sure that the probability estimate that a bit is a 0 added to the probability estimate that a bit is a 1 sums to 1. This is shown in equation (44), (45), (46), (47).

$$q_{ij}^{l}(1) = k_{ij} \cdot P_i \cdot \prod_{j' \in C_{i\backslash j}} r_{j'i}^{l-1}(1) \tag{44}$$

$$q_{ij}^{l}(0) = k_{ij} \cdot (1-P)_i \cdot \prod_{j' \in C_{i\backslash j}} r_{j'i}^{l-1}(0) \tag{45}$$

$$q_{ij}^{l}(0) + q_{ij}^{l}(1) = 1 \tag{46}$$

$$k_{ij} = \frac{1}{P_i \cdot \prod_{j'} r_{j'i}^{l-1}(1) + (1-P)_i \cdot \prod_{j'} r_{j'i}^{l-1}(0)}$$

(47)

In equations (44) and (45) C_i is the set of all check nodes that are connected to variable node c_i. Turning to the Tanner graph shown in Fig. 18, C_{500}={28,74,139,212,228}. $C_{i\backslash j}$ is the set of all check nodes connected to variable node c_i with the exception of f_j.

Once the check nodes have received query messages from the variable nodes at round l, then we are ready for another iteration, l+1. Before moving on to the next iteration we will stop and check the syndrome of our best guess at this point.

$$Q_i^l(0) = k_i \cdot (1 - P_i) \cdot \prod_{j \in C_i} r_{ji}^{l-1}(0) \tag{48}$$

$$Q_i^l(1) = k_i \cdot P_i \cdot \prod_{j \in C_i} r_{ji}^{l-1}(1) \tag{49}$$

In equation (48) and (49) $Q_i^l(0)$ is the best guess at the probability of the bit at variable node c_i at iteration l. If $Q_i^l(0)$ is greater than $Q_i^l(1)$ then our best guess is that the bit at variable node c_i is a 0.

Converting to the log domain turns multiplications into additions when calculating the query messages to be sent from the variable nodes to the check nodes.

$$l_i = \ln\left(\frac{1 - P_i}{P_i}\right) \tag{50}$$

$$q_{ij}^l = \ln\left(\frac{k_{ij} \cdot (1 - P)_i \cdot \prod_{j' \in C_{iJay}} r_{j'i}^{l-1}(0)}{k_{ij} \cdot P_i \cdot \prod_{j' \in C_{iJay}} r_{j'i}^{l-1}(1)}\right) \tag{51}$$

$$q_{ij}^l = l_i + \sum_{j' \in C_{j|i}} \ln\left(\frac{r_{j'i}^{l-1}(0)}{r_{j'i}^{l-1}(1)}\right) \tag{52}$$

$$q_{ij}^l = l_i + \sum_{j' \in C_{j|i}} \mathrm{r}_{\mathrm{j'i}}^{\mathrm{l-1}} \tag{53}$$

And the best guess at iteration l in the log domain is given in the following equation.

$$l_i^l = l_i^0 + \sum_{j \in C_j} r_{ji}^{l-1} \tag{54}$$

<u>Downstream LDPC Code word</u>

The downstream LDPC code is 16,200 bits long with 14,400 parity bits and 1,800 check equations. The forward error correction scheme includes many elements which include bit interleaving, frequency and time domain interleaving and an outer concatenated code. The number of input bits to the outer code is 14,232 as listed in Table 2. So, the overall code rate for the downstream concatenated LDPC code is 14232/16200=87.85%. The method for creating the parity generation matrix, P in equation (2), is described in reference [9].

<u>Upstream Short Size LPDC Code word</u>

The upstream is an OFDMA system, meaning that individual cable modems can transmit simultaneously with each using a subset of the available sub-carriers. In a multi-point to point system, cable modems cannot pool their upstream data to form a large codeword, each cable modem must form a code \word with its own data to transmit. For this reason, a small codeword size is needed on the upstream in order to efficiently transmit small data transmissions. Since a shorter codeword does not have as much coding gain as a longer codeword, a higher coding rate is chosen.

The short codeword length for upstream transmission is 1120 bits with a ¾ code rate. The number of parity bits is 280, and the number of information bits in a code word is 840 bits.

An illustration of the sparse matrix nature and the QC-LDPC structure of the upstream short codeword is shown in Fig. 22.

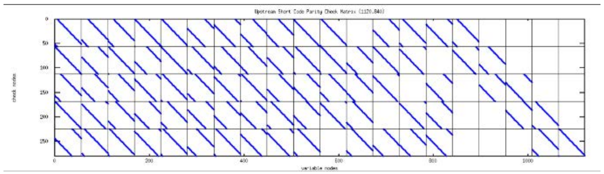

Fig. 22 Spy Diagram of 1120 Bit Short Upstream LDPC Parity Check Matrix.

Upstream Medium Size LDPC Code

The medium size upstream LDPC code word is 5940 bits with 900 parity checks and 5040 information bits with parity check matrix shown in Fig. 23.

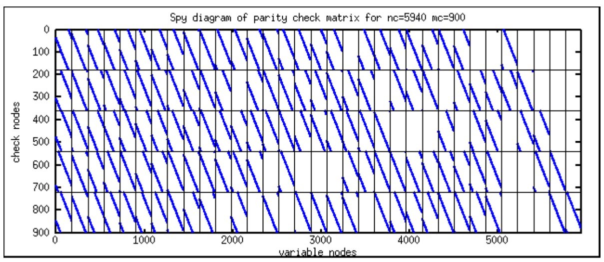

Fig. 23 Parity Check Matrix 5940 Bit Code.

Upstream Long Size LDPC Code

The upstream code word size for long packets is 16,200 bits with 1800 parity bits and14,400 information bits. The parity check matrix is built from a 5 rows by 45 columns base matrix with a sub-matrix size of 360. The parity check matrix is a sparse matrix as shown in Fig. 4.

The base matrices are still being worked on in standards bodies. An excellent description of the LDPC code word formulation using base matrices is found in reference [13] along with example base matrices.

Performance Analysis

Four signals have been described; the downstream signal, and the short, medium, and long upstream signals. All use OFDM and each uses a different LDPC code. Each signal can use various modulation rates. The highest modulation rate analyzed in this paper for the downstream signal was 4K-QAM and for the upstream signal 1K-QAM. Computer simulations were run using Octave on a Linux computer to verify the AWGN SNR threshold and analyze the performance in the presence of noise sources taken from field measurements. The downstream signal will operate error free with an SNR of 36 dB in an AWGN channel. The upstream signals will all operate error free with an SNR of 30 dB in an AWGN.

The cable plant can pick up impulse noise that is short in time duration with wide spectral components, and ingress noise that is long in time duration, yet having narrow spectral width. Impulse noise tends to be created by electrical devices while ingress noise tends to come from radio signals. Ingress and impulse noise are not AWGN since the noise is not white and it can come and go over time. By white noise it is meant that the power spectral density is flat over frequency.

The results presented in this section concentrate on the ability of an OFDM signal with LDPC modulation to operate in the presence of measured impulse noise. Impulse noise of high amplitude is known to be readily present in the upstream, particularly below 20 MHz. An upstream spectrum plot is shown in Fig. 24. This plot shows three DOCSIS 3.0 upstream signals, each 64-QAM modulation with a symbol rate of 5.12 MHz. The blue trace is a peak hold measurement while the black curve is a single sweep. Notice that the SNR is 40 dB at 30 MHz, but is only 15 dB at 5 MHz for the peak hold curve.

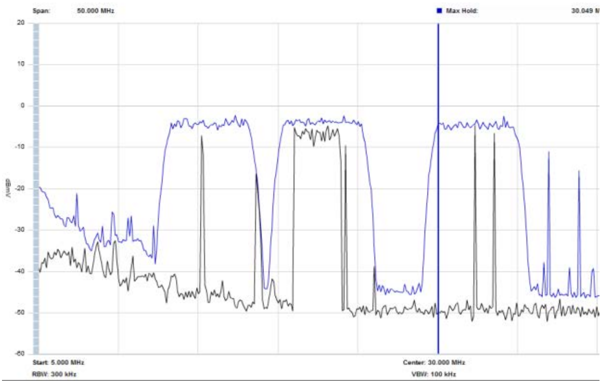

Fig. 24 Upstream Spectrum Plot.

Fig. 25 shows a measured upstream impulse response in the time domain (bottom plot) and the frequency domain (top plot). Superimposed on the spectrum plot is the FFT of a portion of the time domain measurement that included a DOCSIS upstream transmission. Notice that the FFT of the measured upstream impulse closely matches the general shape of the field spectrum plot, with amplitude decreasing with frequency.

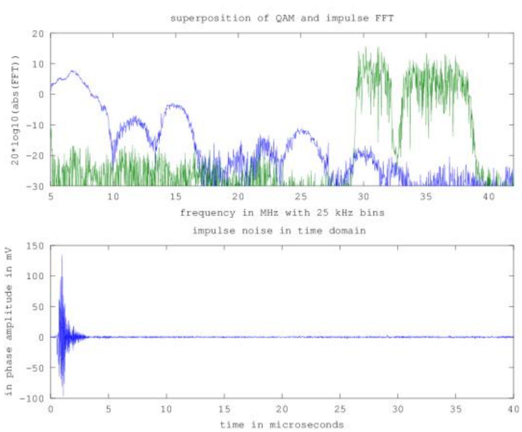

Fig. 25 Measured Upstream Impulse Noise.

Based on the observations made in Fig. 24 and Fig. 25, the upstream bit loading scheme shown in Fig. 26 was derived. The subcarriers fully occupy the 5-42 MHz upstream spectrum. The number of bits per OFDM symbol is calculate in Table 3 and the data rate is calculated in Table 4. With the chosen subcarrier bit loading scheme the data rate in the 5-42 MHz channel is 288 Mbps.

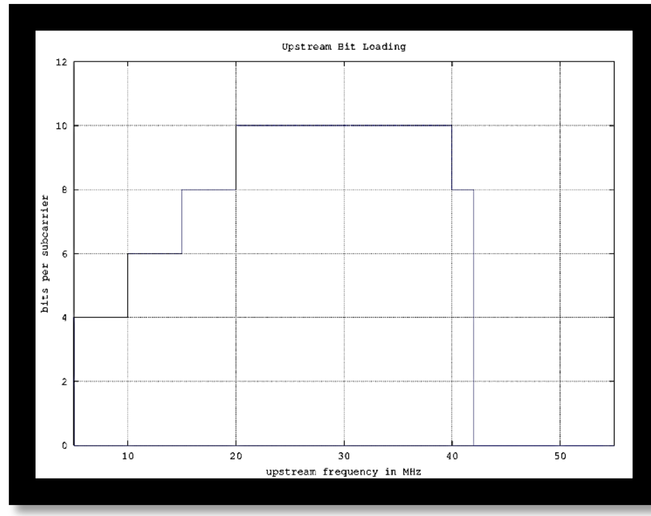

Fig. 26. Upstream Bit Loading 5-42 MHz.

Region	fstart, Hz	fstop, Hz	fspan,Hz	Na	Nbpsc	Nbpr
1	5,000,000	10,000,000	5,000,000	200	4	800
2	10,000,000	15,000,000	5,000,000	200	6	1200
3	15,000,000	20,000,000	5,000,000	200	8	1600
4	20,000,000	40,000,000	20,000,000	800	10	8000
5	40,000,000	42,000,000	2,000,000	80	8	640
						12,240

Table 3. Upstream Bit Loading 5-42 MHz

Useful Symbol time	4.00E-05	s
CP	2.50E-06	s
total symbol time	4.25E-05	s
bits per symbol	12,240	bits/symbol
data rate	288,000,000	bits/s
codeword length	16,200	bits
codewords per symbol	0.75555556	

Table 4. Data Rate with Upstream Bit Loading.

The OFDM subcarriers fill the 5-42 MHz frequency band using variable bit loading, based upon observed noise spectral density. A spectrum analyzer plot of the upstream signal is shown in Fig. 27, measured with an Agilent VSA 89601.

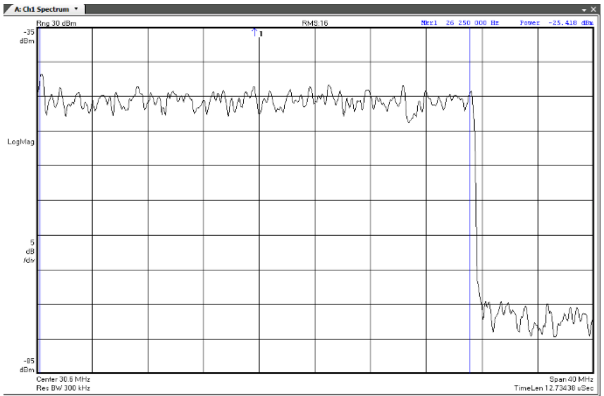

Fig. 27. Spectrum Analyzer Measurement of Upstream 5-42 MHz OFDM signal.

The peak to average power ratio is often of concern for OFDM signals since with so many subcarriers there can be times when the voltage waveform of many subcarriers add in phase, creating a high amplitude peak envelope composite signal. A measurement of the complementary cumulative distribution function (CCDF) of the 5-42 MHz OFDM signal is shown in Fig. 28. A comparison measurement of 4 DOCSIS 3.0 upstream carriers, each with 64-QAM modulation, ATDMA and 5.12 MHz symbol rate [12] is shown in Fig. 29.

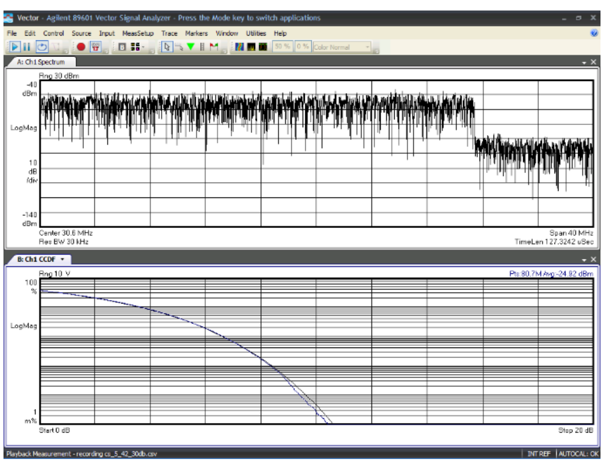

Fig. 28. CCDF Measurement of OFDM.

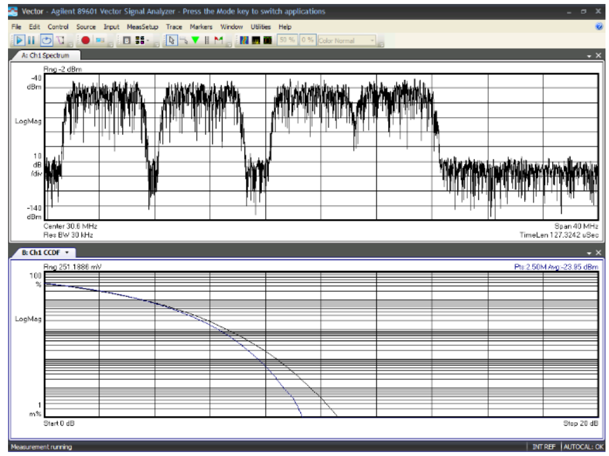

Fig. 29. CCDF Measurement of 4 DOCSIS Upstream Carriers 64-QAM, 5.12 MHz Symbol Rate.

The CCDF measurement of the 5-42 MHz OFDM signal shows that for 99.8% of the time the OFDM signal power is not more than 8 dB above the average signal power. The CCDF of the OFDM signal closely follows the Gaussian distribution. This observation is a direct result of the central limit theorem of probability [7].

Comparing the CCDF of the OFDM signal to the four DOCSIS ATDMA signals, the difference is peak signal power is 0.3 dB at the 99.8% probability point. The OFDM peak to average power ratio is slightly higher than four bonded DOCSIS ATDMA upstream signals.

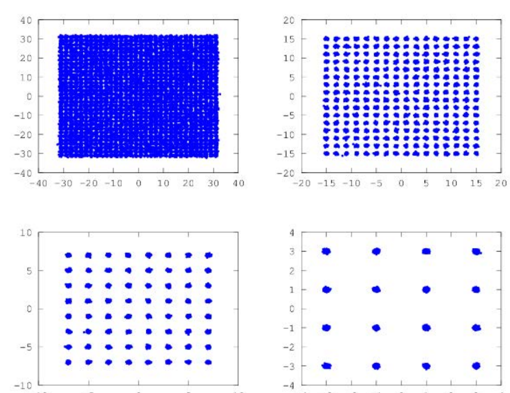

Fig. 30. Constellation Diagrams for the Four Levels of Bit Loading with a Flat 33 dB AWGN SNR.

The constellation diagram of the four types of modulation used in the upstream OFDM signal with a 33 dB SNR is shown in Fig. 30. The impulse noise shown in Fig. 25 was added to the OFDM signal, and the resulting constellation diagram for the 16-QAM modulated subcarriers in the 5-10 MHz range is shown in Fig. 31. The impact of the impulse noise in the 5-10 MHz band is readily evident.

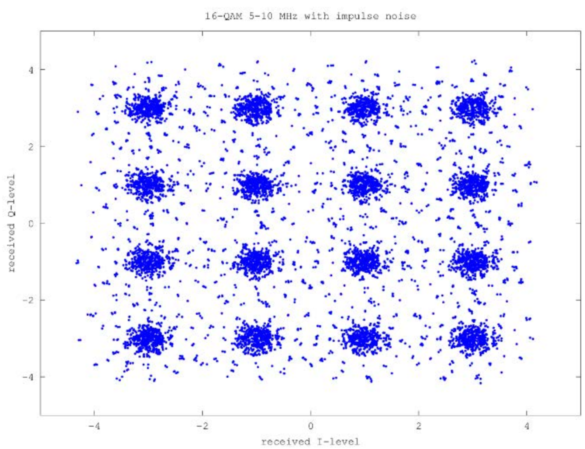

Fig. 31. Impulse Noise Distorts 16-QAM 5-10 MHz.

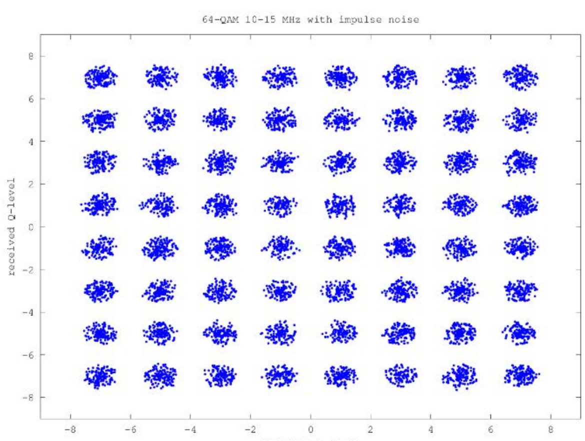

Fig. 32. Impulse Noise Impact on 64-QAM Subcarriers in the 10-15 MHz Band.

The constellation diagram of the 64-QAM subcarriers in the 10-15 MHz band with the impulse noise is shown in Fig. 32. Again, the impact of the impulse noise is readily apparent, but it is clear that the constellation points are still well within decision boundaries. The constellation diagram for 256-QAM subcarriers with the impulse noise is shown in Fig. 33. The constellation diagram of the 1K-QAM subcarriers in the 20-40 MHz band with the impulse noise is shown in Fig. 34. Clearly for the 1K-QAM subcarriers, LDPC iterative decoding will have to work its magic.

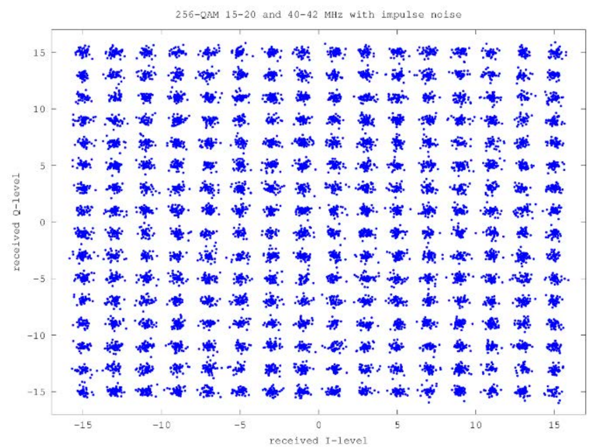

Fig. 33. Impulse Noise Impact of 256-QAM.

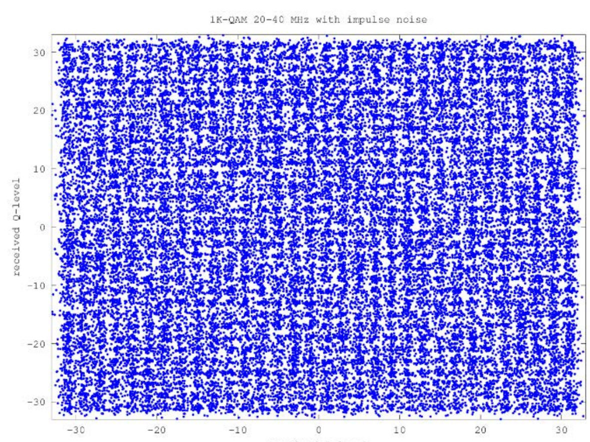

Fig.34. Impulse Noise Impact of 1024-QAM Subcarriers in the 20-40 MHz Band.

The LDPC code used for the analysis of 5-42 MHz upstream OFDM transmission with the field measured impulse noise was the long upstream LDPC code. The codeword length is 16200 with 14400 information bits and 1800 check nodes. The useful symbol time was 40 µs and the cyclic prefix was 2.5 µs. The FFT size was 2048 and the cyclix prefix was 128 time samples. The sampling frequency was 51.8 MHz. The center frequency of the IFFT output was chosen to be at 30.6 MHz so that the lowest frequency FFT subcarrier started at 5 MHz. The decoder used equations (43), (53), and (54), which is a standard flooding sum product message passing algorithm in the log domain.

The general method used in the analysis was to first create 34 codewords, each 16200 bits, using the parity generation matrix. The complex symbols were created by dividing up the codeword bits to match the subcarrier bit loading and mapping bits to symbols for each modulation scheme and normalizing. Then, a signal could be generated by assigning complex symbols to OFDM subcarriers and applying an IFFT to create the time domain waveform. Then, the cyclic prefix could be added along with AWGN noise and delay echo.

The field measured impulse noise was added to the time domain OFDM signal. The amplitude of the impulse was adjusted to match the power spectral density relationship between the field impulse measurement and the DOCSIS carriers. The impulse time domain signal was multiplied by the complex exponential of the center frequency difference in the field measurement and the simulation signal. This aligned the two signals in the freuqency domain. The impulse was adjusted so that the peak of the impulse fell in the midddle of the useful symbol time. The impulse was repeated for every OFDM symbol. In some cases the impulse repeats at a slow enough rate to allow for a time interleaver to correct for the errors due to the impulse. However, in this analysis the impulse repeats every symbol so that a time interleaver would not help.

Then, the log-likelihood ratios for each of the codeword bits were determined. Finally, the log domain message passing algorithm was applied to decode the bits.

Fig. 35 shows a plot of the number of iterations required for the message passing algorithm to eliminate all bit errors in the codeword versus the number of original bit errors in the codeword after demodulation. When the number of bit errors in the 16200 bit codeword was less than 100, no more than 6 iterations were required to get to zero bit errors and a zero syndrome vector. The most bit errors in a codeword was 251, and the highest number of iterations to eliminate all errors was 14. While it would seem that the decoder is not being pushed to its maximum with such a small number of iterations, it was found that increasing the amplitude of the impulse from this point would result in uncorrectable codewords. So while it could be shown that codewords could be corrected by using 30 iterations, when errors were this bad, inevitably some codewords could not be corrected. When the decoder failed to reduce the errors in each iteration, the subsequent iterations actually could increase the number of errors substantially and more iterations made things worse.

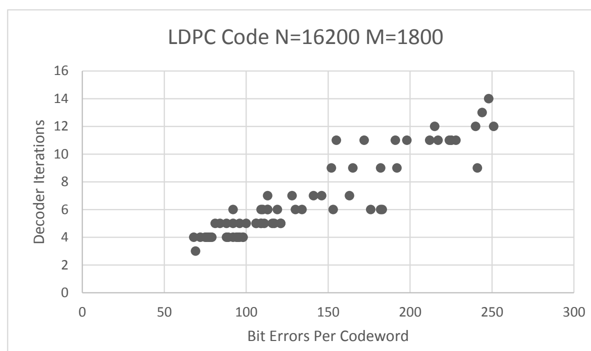

Fig. 35. Decoder Iterations vs. Bit Errors per Codeword for 5-42 MHz Upstream Signal with Impulse Noise.

CONCLUSION

4K-QAM for downstream and 1K-QAM for upstream require signal to noise ratios that are practical to achieve in a cable plant. At 12 bits per symbol 4K-QAM provides very high spectral density, the number of bits per second that we can get out of a Hz of spectrum.

Even at high spectral efficiency, Gbps data rates require wide channel widths. 192 MHz of channel width is required to get 1.87 Gbps of data rate using the downstream 4K-QAM and forward error correction. OFDM is very helpful in making such a wide channel width simple and practical. The FFT algorithm used to create 8,192 subcarriers is computationally very efficient. Using a 204.8 MHz sampling rate and an 8,192 point FFT, the channel is divided into subcarriers that are spaced apart in frequency by 25 kHz. The useful symbol time is 40 µs, which allows for a cyclic prefix of 2.5 µs to be added without sacrificing too much spectral efficiency. As long as reflections in the cable plant have a time delay less than the cyclic prefix they will not contribute to intersymbol interference. OFDM offers a simple mechanism to eliminate intersymbol interference due to reflections. The long symbol time and narrowband subcarriers can average out impulse noise while narrowband ingress noise only impacts a small number of subcarriers.

Finally, LDPC codes have been shown to work well in AWGN channels and also to work well with impulse and ingress noise. The key principals of LDPC coding are using a large parity check matrix that is sparse, thus keeping the computational complexity manageable and still using a large enough sample of bits to have statistically significant information from the channel demodulator. The message passing decoder receives the LLRs for each bit of the codeword. Then, messages representing the probability of each bit are passed from variable nodes to check nodes and back until the codeword is found.

Analysis was done by adding field measurements of noise to the OFDM signals. In many cases, such as in the example presented, the LDPC decoder could easily correct for the errors resulting from the impulse noise. Many field measurements were analyzed in addition to the impulse shown in the example. The example impulse shown is very typical and almost all impulses observed had similar characteristics. The main factor in determining whether or not the decoder could correct all bit errors was the amplitude and frequency response of the impulse noise. At some point as the amplitude of the impulse was increased the decoder failed to correct bit errors. Once the decoder fails to keep reducing errors on every iteration, then the errors start to increase and additional iterations just make it worse. In the example shown, the amplitude of the impulse was adjusted to the highest possible for bit error free decoding. In most cases the amplitude and frequency response of the field measured impulse noise could easily be corrected by the LDPC decoder using the chosen bit loading scheme. However, this was not the case for all field measurements. In cases where the amplitude of the impulse noise or the frequency response of the impulse noise results in unacceptable bit errors, then there are still a number of remedies. One is time interleaving, which can help for a single impulse that does not repeat many times throughout the time interleaving process. Another remedy is to adjust the bit loading, reducing the bit loading to subcarriers most impacted by the impulse.

Still, in general, it has been shown that high modulation rates such as 4K-QAM and 1K-QAM with OFDM and LDPC work well in the presence of field measured noise. Data rates of 1.87 Gbps in a 192 MHz downstream channel width and 288 Mbps in the upstream 5-42 MHz are possible using these techniques.

REFERENCES

[1] R. Gallager, "Low density parity-check codes," IRE Transactions on Information Theory, January 1962.
[2] R.M. Tanner, "A Recursive Approach to Low Complexity Codes," IEEE Transactions

on Information Theory, Vol. IT-27, NO. 5, September 1981.
[3] Tuan Ta, "A Tutorial on Low Density Parity-Check Codes", The University of Texas at Austin.
[4] T.J. Richardson, R. Urbanke, D. Forney, S.Y. Chung, "On the Design of Low-Density Parity-Check Codes within 0.0045 dB of the Shannon Limit", IEEE Communications Letters, Vol. 5, NO. 2, February 2001.
[5] Athanasios Papoulis, "Probability, Random Variables, and Stochastic Processes", McGraw-Hill, 1984, pp. 11, equation (1-8).
[6] Tom Richardson, Rudiger Urbanke, "Modern Coding Theory", Cambridge University Press, 2009.
[7] Proakis J.G., Digital Communications, 2nd ed, McGraw-Hill, ISBN: 0-07-050937-9, 1989.
[8] http://www.tvhistory.tv/1930-ATT-BELL-pg26-27.JPG
[9] Digital Video Broadcasting (DVB); Frame structure channel coding and modulation for a second generation digital transmission system for cable systems (DVB-C2), Draft ETSI EN 302 769 V1.0.2, 2009
[10] MacKay, D.J., Neal, R. M.: Near Shannon limit performance of low density parity check codes, Electronics Lett., Mar. 1997, vol. 33, no.6, pp. 457-458
[11] Research for Development of Future Interactive Generations of Hybrid Fiber Coax Networks (ReDeSign), www.ict-redesign.eu
[12] Data Over Cable Service Interface Specifications, DOCSIS3.0, Physical Layer Specification, CM-SP-PHYv3.0-I08-090121, Cable Labs, January, 21, 2009.
[13] Prodan, R. , Shen, B.Z., Lee, T. "DOWNSTREAM FEC PROPOSAL FOR IEEE 802.3 EPOC", 3/11/2013 http://www.ieee802.org/3/bn/public/mar13/prodan_3bn_01_0313.pdf

APPLICATIONS OF BIG DATA ANALYTICS TO IDENTIFY NEW REVENUE STREAMS & IMPROVE CUSTOMER EXPERIENCE

Anukool Lakhina and Brennen Lynch
Guavus, Inc.

Abstract

Cable Operators have an unprecedented opportunity to leverage an existing asset (big data) to better understand subscriber behavior, improve quality of experience (QOE) for their subscribers, and generate new revenue streams.

The challenge so far has been to address the rapidly growing rates of data volumes, which are inturn being driven by the explosive rise of IP content, tablets & smartphones, sensors and M2M technologies. In order to harness this "big data", operators need to more effectively draw out correlations and insights that have genuine actionable impact to customer experience, revenue, and operations, while simultaneously maintaining day-to-day operational requirements. Communication service providers (CSPs) therefore need solutions to utilize existing network assets more efficiently.

This paper demonstrates how CSPs can tangibly leverage their data assets to drive business value with three applications of big data analytics: IP Video/CDN operations, High Speed Data operations, Customer Care & Network operations.

INTRODUCTION

In the hyper-competitive environment of cable operators and service providers, driving down costs, generating new revenue streams, and retaining/attracting customers are key differentiators. Generating more value from the network and subscribers is key to competitive differentiation in an environment that has a projected 40% growth in data demand on the networks vs 5% growth in IT spending per year [12]. A powerful trend of moving to IP-delivered (packet switched) services [2] and the related innovation within network design/data transport are driving this globally.

At the same time, maturity in network instrumentation technologies has given operators a unique opportunity to gain unprecedented and continuous views into how their customers use & interact with their various services. Therefore, in principle operators can leverage these data assets to offer more personalized services for improved customer engagement, monetize data for increased revenues and optimize network performance for an enhanced customer experience.

However, there are technical challenges to overcome in order to effectively harness this network-derived "big data". The sheer scale of this data – in terms of volume (size), velocity (rate at which this data is generated) and variety (the multiple silos that it resides in) makes it difficult to manage using just the traditional business intelligence and data warehouse technologies.

Even just collecting network data requires the ability to pull data from the edge through the backbone network in real time, without taxing mission-critical infrastructure. Beyond data collection, in order to be useful this network-generated data needs to be dynamically fused and correlated with static & reference data sets to produce key causal relations. This refers to a variety of data sources from network-generated to billing, cost & pricing catalogs, subscriber demographics, business

KPIs, and other external datasets. Thus, *a new data fabric* is needed that not only has the capability to ingest huge volumes of distributed network data being generated at extremely high velocity from a variety of elements, but also fuses the network-generated data with business context that exists already in information systems today (e.g., billing plans, subscriber demographics, cost models, etc) to enable discovery with true context.

This environment presents many considerations and challenges to the development of a big data analytics structure. First, it must scale to address the massive volume of daily petabyte-size streaming datasets. Second, the high cardinality - large number of distinct attributes to each dataset - creates additional compute considerations when designing a data fabric that can efficiently process and analyze the raw data. Third, the supporting platform needs to be carrier-grade, meaning that it must adhere to high availability standards of any network-ready system. Fourth, the platform must be able to function in a distributed architecture wherein processing can happen at the edge of the network resulting in reduced transport costs associated with moving petabyte-scale raw data to a central location. All these requirements requires a departure from the traditional model of 'centralized, store-first' batch analytics to a distributed, compute-first continuous analytics model wherein streaming data is continuously processed & analyzed at the edge.

Of course, the ultimate value of this fabric lies in rapidly creating business aware applications that decision makers can use and that trigger automated, closed-loop business processes. And so finally, the platform needs to be an integrated & holistic stack, from data ingestion to processing and finally visualization with APIs to allow for applications to be built rapidly.

The rest of this paper presents examples of some applications for cable operators.

APPLICATIONS OF BIG DATA ANALYTICS

More than integrating mining algorithms, successful data science must be able to view business problems from a data perspective [13]. Once the compute is pushed out to the edge and analytics can be timely, then decisioning applications can be built on top of the big data fabric created to enable real-time, relevant, actionable insights (optimize, delight the customer, and discover new revenue opportunities).

Creating a business-aware decisioning application that sits on top of the described big data fabric is key to unlocking the data's inherent value. A software application driven by a business or network user can be described in a number of contexts to enable actionable insights that optimize, reduce costs, and identify new revenue streams.

Three examples on how a decisioning application can be built per context are provided. Each specific context- IP Video/CDN, Subscriber Analytics on high speed data (HSD) services, and integrated Customer Care/Network Operations- enables discovery of highly impactful revenue-driving and cost saving insights.

IP Video – IPTV, OTT & CDN

Consumer expectations have been moving from plain linear programming to video on demand, HD-quality and interactive viewing experiences. Increasing consumer trends of ditching linear TV for IP-delivered media has led Service cable operators to adapt in a number of ways such as rolling out competitive services. HFC plants are moving to support 10Gbps capacity to handle traffic demand; new media handling and translation capabilities by home gateway devices allow convergence of services to meet (future) consumer expectations [3]. While the environment evolves, the key challenges that face the business and network executives are fundamentally the same: retaining existing customers and attracting new ones. Moreover, with increasingly customizable subscription packages available to the customer, monitoring the state of the business and identifying-prioritizing high value focus areas are key differentiators to the cable operator.

Service providers deploying competitive IPTV services must increasingly deliver IP Video to multiple devices beyond the traditional set top box including computers, mobile phones, tablets,.. at a high QoS to satisfy the consumer's expectation of consistent experience. Since they own the "last mile", service providers are in the unique position of being able to offer efficient IP services deployed at the edge. This player data (generated by the end subscriber) are raw customer interaction metrics- around how subscribers use the network. Fusing this raw data with additional datasets allows the operator to calculate and monetize high-value actionable insights such as propensity to churn through user QoE issues or subscriber profitability. In a business context product engineers are provided transparency towards granular demand and campaign effectiveness. Marketing teams can target high churn risk subscribers for remediation or enable relevant delivery of advertisements to profitably drive advertising revenue.

Service providers with managed CDNs must deliver content with the same high quality. This application data, which may be procured from a third party by the provider in some cases, is generated by the video serving equipment (also known as CDN delivery nodes). The operator can perform network optimization with status and trend information by content or other CDN metrics (location, node health, etc). In addition, content analysis provides avenues towards targeted marketing. As operators deliver more IP content and transition head-end service delivery to CDN architecture [2] it becomes a richer and more relevant perspective for analysis.

Player data and application data each provide a unique measure of analysis that leads to actionable discoveries. Both prime data sources must be fused with the same or similar catalogues, location databases, as well as categorization databases in order to resolve all data to necessary metrics such as subscriber and topology. These two unique measures are complementary and logically related as two parts of one whole application (with the necessary reference/catalog databases)- deployed with input of one, or best, both prime data sources:

IPTV Sources- Player Data

Generated by the end subscriber, player data is mediated and aggregated from a number of raw data sources including guide-based and data generated from the delivery software located at the client player. To gain full perspective fusion must also occur with any geolocation service (enabling ISP identification) and information around subscriber DVRs, entitlements, and purchase activity.

The analytics solution built on player data provides a number of unique insights around traffic trends, subscriber engagement and content metrics.

To reveal subscriber QoE (ie traffic anomalies through consistency & continuity) traffic is resolved to subscriber; then topology to reveal potential drivers. Marketing solutions around churn enable near real-time action on high-risk subscribers. Traffic metrics on broader terms enable dynamic product/network engineering efforts such as variant analysis (how changing one traffic metric - ie delivery - will affect others).

Content discovery to the subscriber or to any part of the network allows behavioral analysis around specific product engineering, as well as targeted marketing and campaign activity. For example, subscribers can be categorized into certain states based on activity, trials (in-trial/ex-trial & activity within trial period), transaction history, subscription package & tenure, etc. The operator can then analyze profitability, demand, success, and identify product engineering decisions driven by data produced by the operator's customers.

Analysis of player data generated by interactions with a Service Provider's IP Video service enables a rich set of actionable insights. An appropriate software application enables executives and other operators to closely monitor, analyze, and identify future trends in the business. It is an open solution that benefits from additional data sources for budget overlay calculations and specific costs. The solution supports Marketing, Product Engineering, Customer Care, Sales, and Financial executives in regular monitoring and course-correction strategies.

CDN Sources- Application Data

Video-serving equipment or CDN delivery nodes generate the application data relevant to content distribution. Application data can be aggregated, processed, and correlated from a number of raw logs or data sources.

The primary source of data is from Delivery Node logs (specific to vendor such as Velocix, Cisco or internal/proprietary logs). The data enables resolution of traffic to subscriber or further granulated into delivery (unicast, ABR, etc). In addition by generating content resolution charts for both constant bit rate (CBR) and adaptive bit rate (ABR) encoded media (requiring associated catalogue data and access logs) the operator can effectively analyze the most efficient and highest-quality method to deliver content. In this manner network engineers can monetize data to implement the most efficient processes and expansions.

Additional data sources to provide a full perspective can include proxy logs from upstream nodes to provide visibility of network utilization from the delivery nodes to the upstream devices; catalog data such as Asset Management System databases; and lastly, Session Manager (created when users select an asset to download) and License Manager logs are used to correlate content accessed to subscribers. Fusion of the above-referenced application data with subscriber/topology-resolution data such as billing/CRM enables near real-time analytics to deliver timely insights on the performance of CDNs.

Content analysis may compare performance of different assets over selectable time periods and allows a marketing or engineering user to: analyze usage patterns, anticipate demand, and provision the CDN efficiently to meet customer requirements at minimal cost. In addition to identifying popular assets, access information etc, one can view measures around how subscribers access using ABR or HLS content and CBR content to increase efficiency and QoE through intelligent data-driven engineering decisions.

Approaching CDN-generated data from a geographical distribution of client access enables the marketing user to deliver titles on a per-demand (down to geographical market) basis. Network engineers can analyze best-locations for network expansions to capitalize on demand and increase the efficiency of expansion.

Device-level analytics is becoming increasingly important as subscribers access more content on devices (tablets, smartphones, etc) other than their primary television/STB. Product engineers can prioritize support of operating systems based on near real-time demand of content. Classifying popularity amongst devices and content-per-device allows for data-driven marketing to users of popular devices & contents.

Network resource utilization is an important perspective enabled with application data. By quantifying metrics such as specific CPU utilization of nodes/caches, HTTP transactions, cache hit ratios etc, the user is provided transparency towards demand on network resources. This perspective allows network engineers to decide when to expand the capacity of the network, view how the CDN is routing requests, and detect possible problems with delivery nodes and caches.

Selective grouping and intelligent fusion of application data with relevant peripheral datasets enables an extremely powerful solution for the CDN operator or Service Provider who can access the relevant CDN data third party. Delivery of these insights in a drill-down approach allows the operator to capitalize on market opportunities, identify important trends, and identify/prevent problems that could have impacted network performance, customer satisfaction and loyalty.

__IPTV, OTT & CDN Summary__

Application data and player data provide a strong set of complementary analytical insights. Correct fusion of both player data and application data into a structured software application can enable the Service Provider to identify a number of actionable insights from a per-subscriber level for IP Video services and a content/network level for CDN deployments. These actionable

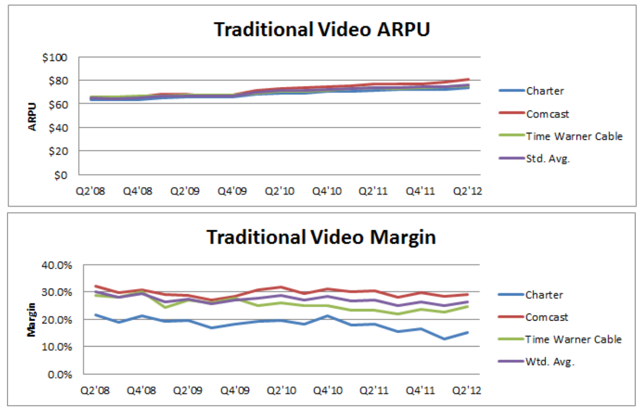

Fig. 2: Video ARPU rises; Video Margin falls. *Representation*: increasing cost of providing acceptable QoS under growing network demand is not equilibrated by ARPU. Net cost of increased bandwidth usage is transferred to the Cable or Broadband operator. *-data from [6]*

insights are numerous, and can be classified with three broad use case fields. Customer experience metrics allows CDN decisioning; viewer measurement enables marketing to perform ad decisioning; and network engineers can leverage asset analytics for network optimization. The application can be structured in closed-loop format to perform necessary network optimizations based on anomalous activity (changing serving CDN, adjusting bitrate, etc) and marketing operations (interfacing with ad decisioning engine).

Subscriber Analytics- High Speed Data

Communication service providers operate in a dynamic and competitive environment of technological advances and elastic consumer demand. Increasingly consumers are moving to IP-connected services, especially in the instance of moving from legacy QAM video to IP-streamed video- generally delivered OTT. Adapting to this changing demand trend and optimizing product/network engineering & deployment are often key differentiators between service providers.

As a result Cable operators and Broadband providers are continuously trying to increase the depth of understanding on how, when, and to what topology level their subscribers access data. They require tools that will increase the efficiency of delivering bandwidth to meet subscriber demand and allow for transparency to congestion/Fair Share/AUP (acceptable use policy) on a per-subscriber (& per-device) level. Most importantly, they require tools that will fuse consumption data with marketing segmentations and demographics, while preserving subscriber privacy. Figure 2 clearly illustrates the operators' collective need to create value from high speed data services with the current trend [6]. Fusion on a real-time, scalable basis creates powerful marketing and product/network engineering applications through calculations such as profitability and retention risk (churn propensity).

Subscriber Analytics- Data Sources

The architecture of a service provider's network allows for extraction of usage-based information through the DOCSIS specification IPDR (IP detail record). IPDRs can be pulled from every CMTS throughout the network and enable per-subscriber discovery (with some peripheral data fusion ie DHCP logs and BSS for subscriber mapping). Key IPDR output details are: traffic as service flows, QoE (ie dropped packets), and MAC address (for cable modem) [15].

Content discovery is enabled through the deployment of deep packet inspection (DPI) probes throughout the network. In an age where most media is ABR streamed and may be coming from multiple CDNs, per-packet analysis is needed to resolve content. As many operators roll out DPI probes in an incremental, per-market fashion, aggregate trends of applications/content may be derived from a small deployment (ie 5-10%). Full deployment of DPI probes enables content discovery per-subscriber down to the domain level, and in some cases, device.

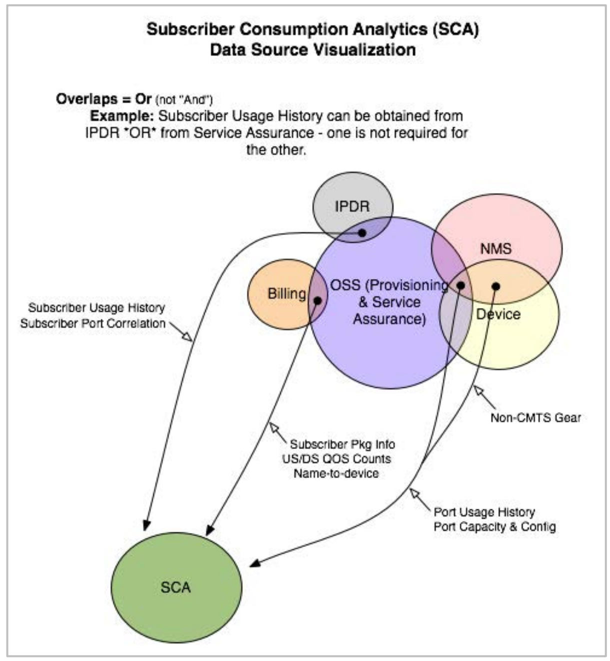

Fig. 3: Subscriber Analytics data sources

Peripheral data such as Billing/BSS and DHCP logs are required for mapping to subscriber. Additional data sources can enrich a solution including marketing-based, OSS/NMS, legacy systems usage, etc. These additional sources can allow for a number of enrichments described below in Use Cases.

Use Cases- Subscriber Usage

Service providers must consistently deliver high-quality, high-speed data to retain subscribers and maintain/increase customer sentiment. Subscribers can experience QoS issues through traffic delivery or specifically through self-congestion (self-imposed QoS due to exceeding traffic delivery limits set by service tier).

Peak time traffic drives operator cost and potential network imbalances with unexpected load. Calculating subscriber profitability from peak time use (amongst a number of other inputs ie care interactions, etc) drives intelligent engineering decisions. Traffic usage can be modeled from the subscriber & topology to ultimately provide visual trend insights towards drivers (specific content, services etc) and also produce subscriber lists for targeted marketing campaigns (ie upsell self-congesting users).

Use Cases- Content & Prospecting

Content analysis is an important source of behavioral information to enable internal product engineering analysis and targeted marketing engagements. Content can be analyzed to the CDN (and further to the geography, subscriber) to discover potential prospecting engagements.

Categorization of content enables powerful prospecting/marketing applications. Defining traffic to a specific category (such as paid-for OTT video traffic or managed hosting services) allows a product engineer to quantify & predict demand and value (relative usage of a particular content vs all others). A business user can then identify appropriate subscribers for a targeted marketing engagement (ie rolling out internal service).

Use Cases- Scorecard

When creating a solution around HSD usage with marketing & product engineering perspectives, dynamic customer segmentation is key. Segmentation allows the user to identify common attributes amongst subscribers that access certain content, participate in certain traffic usage patterns, etc. This can be defined by demographic information when available.

Segmenting subscribers into actionable groups allows the business user to mitigate churn and increase profitability (on the network, products, subscribers). Important calculations can be used to segment subscribers based on their traffic and content usage, (as well as on subscriber attributes). Calculated scores used to segment subscribers may include:

- Profitability: per-subscriber profitability based on inputs such as usage, peak time, care interactions, revenue-consumption, etc
- Upsell: propensity to upsell based on network interactions, self-congestion & QoE, etc
- Retention Risk: calculates per-subscriber churn propensity. Can be defined by a number of inputs such as OTT spend, consumption,..
- Agony: defines pure QoE to subscriber to identify high-impacted subs

Example use of these scores would include Retention score for targeted marketing, Agony score for special treatment (engage/elevate subscriber with service reps), Profitability score for product engineering decisions, and general segmentation for campaign management.

Subscriber Analytics Summary

The ability to leverage data around delivery of HSD will provide CSPs a significant competitive advantage in reducing relative cost of the increasing bandwidth demand, increasing QoS, staying relevant to customer demand and identifying large-volume uploaders that disrupt revenue balance. CSPs can mine high value from their data with a software application that combines usage, topology, content, customer segmentations and score calculations. Integrated systems with Care can export high churn risk subscribers as they trend upwards, enabling closed-loop business processes. Such a holistic solution will enable CSPs to adapt to the trend of OTT services; rolling out competitive services and identifying key operational changes in the network.

CSPs with the opportunity to identify and analyze, in near real-time, the potentially impacted customers. Ultimately software applications can be structured closed-loop to enable true machine-to-machine (M2M)

Customer Type	ARPU	Churn	CLV
Analog	$45.00	2.50%	$222
Digital	$67.00	3.00%	$599
Video/Data	$100.50	2.00%	$2057
Video/Phone	$93.50	2.00%	$2062
Data/Phone	$71.00	2.00%	$1790
Triple Play	$133	1.00%	$5642
Triple Play	$133	2.00%	$2972

Fig. 4: Results of Customer Lifetime Value (CLV) Formula: CLV = Present value (ARPU – COGS – Care Costs). Maximizing CLV = Minimizing Care cost. *Source: [14]*

software collaborations (with internal CSP systems) that can also enable necessary actions such as notifications, IVR call deflections, and targeted truck roll modifications. Allowing for automated business processes based on the data maximizes efficiency in reducing care interactions associated with network events.

Network Impact on Care Interactions

In an environment that generally experiences high churn rates, the quality of service provided is often the key differentiator between companies. With that, CSP's are continuously trying to improve the customer experience in order to retain existing customers and attract new ones. Further, CSP's cost of customer care significantly impacts their profitability and thus requires in-depth insight into cost drivers. More importantly, they need tools to increase the efficiency of delivering a higher level of care.

Currently CSPs have limited insight into the correlation between network events (NE) and care events (CE). For example, a multitude of NEs can trigger spikes in call volume, unnecessary CSR engagement and costly truck rolls. Correlation of NEs and CEs can provide a holistic understanding of key call volume drivers, whether they are a result of anomalous NEs, or point to deeply patterned issues such as device interoperability or faulty equipment. Scalably-built correlations provide

Subscriber churn is a large concern for any CSP. Churn is directly tied to customer experience; a poor customer experience is directly attributable to a larger churn rate. Even a very modest reduction in churn will result in a significant increase of ROI (& vice versa). In fact, Dr. Rizutto at the University of Denver financially quantifies customer experience with his Customer Lifetime Value (CLV) formula [14], the results of which are modeled in fig. 4**.** This problem of churn clearly compounds the cost-to-care for CEs.

Data Sources-Network Impact Care

There is a wide variety in the data sources required to gain a holistic view of the network impact on care interactions. They may be categorized as follows:

1. Billing, Customer care/IVR
2. Provisioning, CPE, STB/Guide
3. Network & Maintenance (ie affected CMTS, node,..)
4. Benchmark & KPIs

Creating a relevant and scalable causal relations structure for the above data sources is a required function of the data mining. This enables near real-time root cause analysis & mediation of care-driving issues.

Use Cases- Care Interactions

Currently lengthy manual intervention is generally required to resolve anomalous care activity to the network-based root cause(s). Care agents require transparency towards network operations that might be driving call/ticket volumes, and ultimately need to be empowered to perform near real-time mediation actions. Therefore root cause analysis and correlating associated impacted subscribers are key application features. This can occur through a process "Relative Commonalities", discussed in detail below.

Carriers are currently blind to drivers of calls that do not result in a ticket (ie simple modem reboot solves the problem). With no other metadata attached other than a call record it is extremely difficult to investigate a potential root cause. Using an application method to correlate subscribers to topology and other contexts (ie new provisioning file or other change management)- a method described below as relative commonalities- a data-driven application can provide transparency towards these obscure yet key call volume drivers.

Use Cases- Problem & Fix Codes

Customer service representatives (CSR) assign problem codes and fix codes to define subscriber (& network) troubles. Analysis by these metrics allows insight into persistent, care-driving trends and investigation of CSRs resolution path.

For example, root cause analysis for repeat interactions tied to obscure issues can be enabled by looking at the trend of problem code assignments. Once a repeat trend is identified with the correct fix isolated, all subscribers pertaining to incorrect ticket assignments can be appropriately assigned the correct fix to eliminate associated care interactions.

Use Cases- Network Events

IT teams - maintenance and network event ticket owners require a tool to perform effective follow-up analysis and furthermore alert them to the presence of anomalous care interactions associated with a network event. Currently, network-generated events such as affected elements or triggered KPIs enables identification of anomalous network events. By correlating care interactions the operator can easily perform the proper follow-up on maintenance and network events by identifying associated troubles (care interactions). Using network-generated data as the first indicator, an application can be structured to:

– Identify subscribers that are potentially impacted by a Network Event, (scheduled or spontaneous)
– Correlate CRTs with a defined geography, type of service, device, network entity, & network element to find root cause
– Generate a list of the potentially impacted subscribers to enforce action (see fig 5).

Relative Commonalities & RCA

Root cause analysis is an integral part of a solution that will enable not only identification of anomalous events but also discovery of the contributing factor(s). This can be made difficult, as not only network failures but also multiple device interoperability can be the root cause. For example, a particular provisioning file pushed to a specific cable modem (& even on a particular CMTS) may produce interoperability issues and are extremely hard to identify without correlating care interactions, to subscribers, to network topology. This may be achieved with a concept called 'relative commonalities' that removes the ambiguity that would visually arise in the graph when a lower network element (such as the Node) is necessarily affected due to the malfunction in the higher network element (such as the CMTS), and allows the operator to clearly view the impact of each specific net work element without any clarity loss stemming from issues of cascading problems.

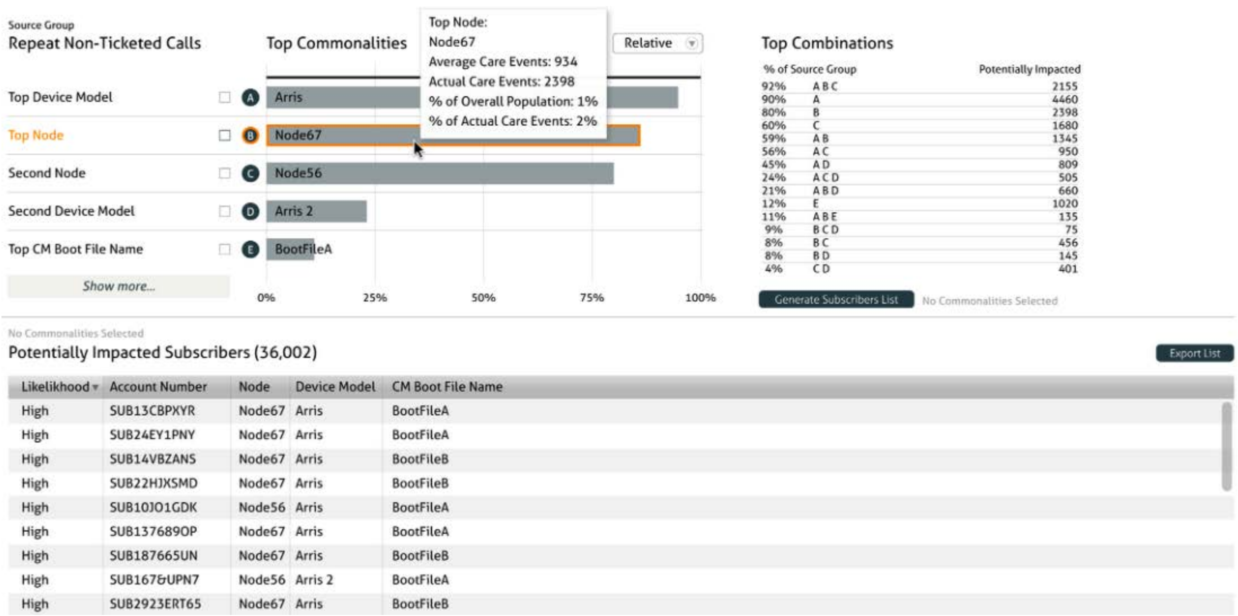

Fig 5: visualization of 'relative commonalities' concept for root cause analysis to produce list of impacted subscribers to be acted upon

powerfully enables identification of one or multiple contributing factors.

Relative commonalities (see Fig 5): this concept of correlating root cause of an anomaly to a network element, customer device, boot file, etc. takes into account network hierarchy, relationships between network elements and the distribution of devices, boot files, etc. throughout the network. Take as example the case of correlating a specific subscriber's troubles to both their faulty Node & their faulty CMTS. In this example, "Node" would be removed from impacting the Relative commonalities weighting, while CMTS would be defined as the problem for that particular subscriber in the aggregation process. In other words, this

Another example of the concept would be taking into account the overall distribution of devices, boot files, etc. If 75% of users have a certain device and 75% of calls are from customers with this same device- the call distribution for devices is the same as for the overall population and that points to some other issue. If however 10% of customers have a certain device and 50% of calls are from customers with this specific device - that would weight as an indication that something is wrong with that device and potentially should be investigated.

Network & Care: Summary

Correlating network & care related data and resolving to subscriber and to network topology creates an extremely powerful analytics application that has high potential to reduce care cost. Reduction of calls, tickets, truck rolls, and MTTU (mean time to understand) issues are high-impact cost reduction benefits. This is enabled through near real-time identification of care and network anomalies to allow discovery of root cause and affected subscribers. Immediate correlation allows the care and IT organizations to interface and immediately perform necessary actions – deflect calls, cancel truck rolls, push a fix, etc. Once this correlation is established, alerts based on templates can be created and even pushed down to markets to enable real-time action on automated anomaly detections. In addition, a true closed-loop system using software collaborations (M2M) with internal ticketing and truck roll systems is extremely efficient in identifying and remediating in near real-time.

CONCLUSION

Creating an end-to-end data analytics solution based on key tenets of scalability and low-latency is key to monetizing carrier-scale data. Data must be relevantly fused with business perspective and modeled to create end-to-end decisioning applications.

In summary, the market outlook for CSPs is undergoing a period of extreme change. The ability to leverage data will provide CSPs a significant competitive advantage in reducing cost to care, minimizing churn, improving customer experience. In addition, CSPs will need to adapt to the trend of OTT services; rolling out competitive services and identifying key operational changes in the network. Scalable software applications with integrated closed loop business processes maximize the operational efficiency of a service provider. Data-driven applications identify key areas to both reduce cost (ie delivering care) and increase revenues (ie targeted marketing). Deploying relevant analytical applications bestows upon CSPs a unique opportunity to gain competitive advantages and maximize profitability.

REFERENCES

[1] Borthakur, Dhruba, Kannan Muthukkaruppan, Karthik Ranganathan, Samuel Rash, Joydeep Sen Sarma, Nicolas Spiegelberg, Dmytro Molkov, Rodrigo Schmidt, Jonathan Gray, Hairong Kuang, and Aravind Menon. *Apache Hadoop Goes Realtime at Facebook*. Tech. N.p.: Facebook, 2011. Print.

[2] Brown, Dave, and Asha Kalyur. *Death of the Headend: How IP Will Transform Cable Services*. Tech. San Jose, CA: Cisco Systems, 2010. Print.

[3] Cheevers, Charles. *Reviewing Cloud-Based Versus Home-Gateway IP Migration Strategies*. Tech. N.p.: SCTE, 2012. Print.

[4] Ciferri, Cristina, Ricardo Ciferri, Leticia Gómez, Markus Schneider, Alejandro Vaisman, and Esteban Zimányi. "Cube Algebra: A Generic User-Centric Model and Query Language for OLAP Cubes." *International Journal of Data Warehousing and Mining* (2012): n. pag. Print.

[5] Huang, Yiyi, Nick Feamster, Anukool Lakhina, and Jim Xu. "Diagnosing Network Disruptions with Network-wide Analysis." *ACM SIGMETRICS Performance Evaluation Review* 35.1 (2007): 61. Print.

[6] *Internet Video, December 2010*. Rep. N.p.: SNL Kagan, 2010. Print.

[7] Lakhina, Anukool., John W. Byers, Mark Crovella, and Ibrahim Matta. "On the Geographic Location of Internet Resources." *IEEE Journal on Selected Areas in Communications* 21.6 (2003): 934-48. Print.

[8] Lakhina, Anukool, Konstantina Papagiannaki, Mark Crovella, Christophe Diot, Eric D. Kolaczyk, and Nina Taft. "Structural Analysis of Network Traffic Flows." *ACM SIGMETRICS Performance Evaluation Review* 32.1 (2004): 61. Print.

[9] Lakhina, Anukool, Mark Crovella, and Christophe Diot. "Diagnosing Network-wide Traffic Anomalies." *ACM SIGCOMM Computer Communication Review* 34.4 (2004): 219. Print.

[10] Lakhina, Anukool, Mark Crovella, and Christophe Diot. "Mining Anomalies Using Traffic Feature Distributions." *ACM SIGCOMM Computer Communication Review* 35.4 (2005): 217. Print.

[11] Lu, Jiamin, and Ralf Guting. *Simple and Efficient Coupling of Hadoop With a Database Engine*. Tech. N.p.: FernUniversität in Hagen, 2012. Print.

[12] Manyika, James, Michael Chui, Brad Brown, Jacques Bughin, Richard Dobbs, Charles Roxburg, and Angela Byers. *Big Data: The next Frontier for Innovation, Competition, and Productivity*. Rep. N.p.: McKinsey Global Institute, 2011. Print.

[13] Provost, Foster, and Tom Fawcett. "Data Science and Its Relationship to Big Data and Data-Driven." *Mary Ann Liebert* 1.1 (2013): 51-59. Print. -

[14] Rizzuto, Ron. "Customer Lifetime Value" NCTA Cable Show, May 2012, Boston, Massachusetts.

[15] Sundelin, Andrew. *What Counts? Accurately Accounting for End-User Traffic with IPDR*. Tech. Boulder, CO: Applied Broadband, 2010. Print.

KEY ACRONYMS

- ABR: adaptive bit rate encoding
- ARPU: average revenue per user
- AUP: acceptable use policy
- CBR: constant bit rate encoding
- CE: care event
- CMTS: cable modem termination system
- CDN: content delivery node
- CLV: customer lifetime value
- CPE: consumer provisioning equip.
- CSP: communication service provider
- CSR: customer service representative
- DOCSIS: data over cable service interface specification
- DPI: deep packet inspection
- HDFS: Hadoop distributed file system
- HLS: HTTP live streaming
- HSD: high speed data
- IP: internet protocol; packet-switched
- IPDR: IP detail record
- IVR: interactive voice response
- KPI: key performance indicator
- M2M: machine-to-machine
- MTTU: mean time to understand
- NE: network event

- OLAP: online analytical processing (context of cube computing)
- OSS/BSS: operating/business service software
- OTT: over-the-top; cord-cutting
- PCA: principle component analysis
- QoE/S: quality of experience/service
- RCA: root cause analysis
- STB: set-top box

Applying the Software Defined Networking Paradigm to MSO Commercial networks

Nagesh Nandiraju , Ph.D. and Sebnem Ozer, Ph.D.
Arris Inc.

Abstract

Today, a growing number of enterprises and SMBs are adopting cloud infrastructure for data storage, computation and services, which is driving increased traffic in unpredictable patterns. Enterprises are seeking greater flexibility from the access network infrastructure to effectively address demands, such as dynamic bandwidth and ad hoc services, from these new cloud services.

To successfully compete and lead in this market, MSOs need to evolve their current infrastructure into a more agile network that is capable of launching new services with much shorter lead times. An overarching common control and management layer is required to address the needs in a timely and cost effective way.

Software defined networking (SDN) presents an interesting concept of decoupling the control and data planes. As traditional, and even more modern, network architectures struggle to cope with dynamic applications and services, applying SDN paradigms will lend towards an effective solution to solve the operators' challenges. This paper highlights how SDN principles can be operationalized to result in a network infrastructure platform that is easy to manage and also presents new opportunities to monetize from the network.

INTRODUCTION

Over the past few years Multiple Service Operators (MSOs) have made successful strides in the business services segment by competing against incumbent players. Just as they are feeling comfortable by optimizing their in-house operations teams for provisioning and delivering the services, new challenges are arising.

With the rise of data centers and maturity of services over cloud infrastructure, a growing number of enterprises and SMBs are migrating towards a public or private cloud infrastructure, consolidating their server rooms (see Figure 1). This presents a change in the network traffic patterns where its not only a North-South traffic (Server-Client) but a lot of East-West traffic between sites or branches and between branches and data centers.

To successfully lead in this market, MSOs can present a differentiating point by offering a seamless interface for the customers to dynamically control the bandwidth needle of their network pipes. Furthermore, while some MSOs have been deploying, others are considering new business models such as mobile backhaul, managed services and wholesale service models. They are also expanding their services beyond SMBs to serve large enterprise customers. Larger enterprises generally have multiple business locations geographically dispersed and span across different service providers' footprint.

In this paper, we will first describe the objectives and requirements for the new dynamic business services from the enterprise customer's and service provider's point of view. We then present a couple of approaches that can be used to achieve the desired flexibility. One way to address these new service models would be by integrating the DOCSIS backoffice with the core network provisioning tools and publishing the noth-bound APIs that 3rd party or in-house software tools can utilize to present a friendly interface to the customers. Another approach is to adopt the SDN paradigm. SDN presents an attractive concept of being able to virtualize the network and can control, configure and program the network (and all network elements) as and when needed. A deeper look into SDN and methods of adopting SDN principles can prepare the MSOs to transform today's network into a future network that is very agile, extremely easy to manage and capable of launching new services with much shorter lead times. As networks continue to evolve, we will provide some insights into how hybrid approaches and adopting SDN incrementally will ease the transformation of today's networks into more flexible and capable networks.

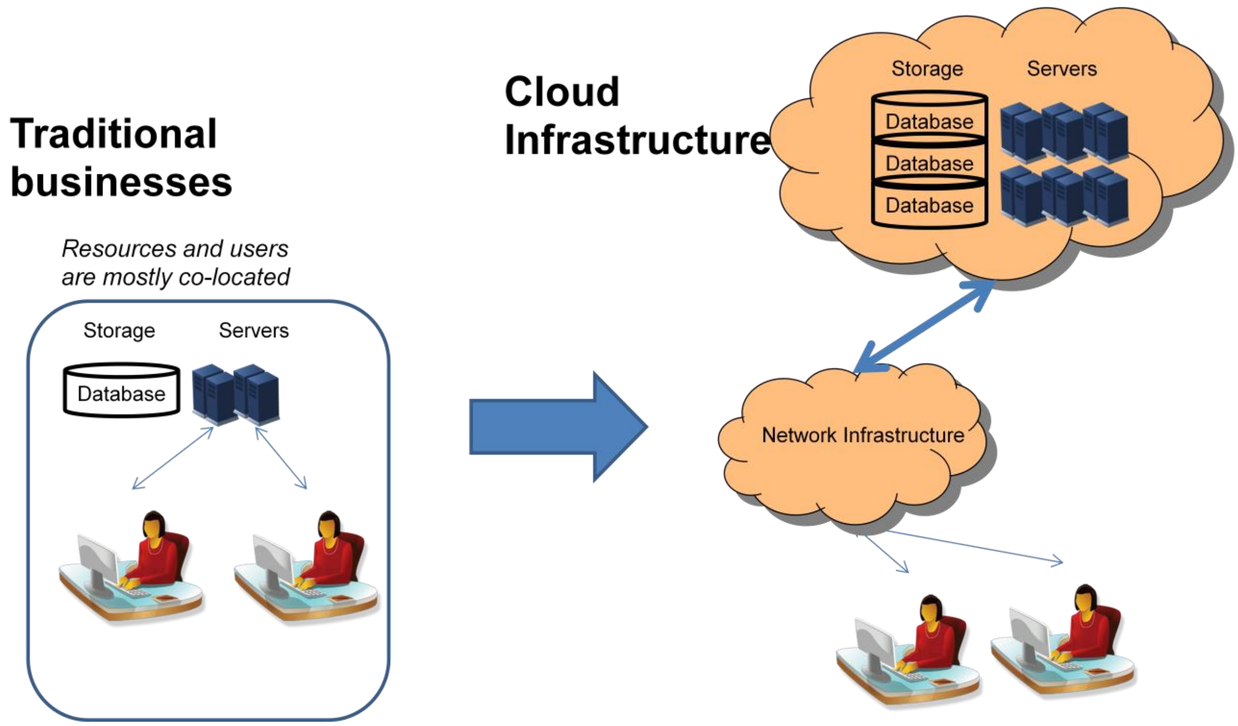

Figure 1. Enterprise Infrastructures migrating to Cloud

BUSINESS SERVICES BY MSOs

Cable operators have been providing residential services for many decades while their business services are relatively new and their current market share is smaller but steadily increasing. Cable operators started to sell their residential triple plays services to small to medium business customers, mostly companies with one location and few employees, local government, healthcare and education. However, increasing revenues and even bigger market opportunities have been a big motivation for cable operators to extend their business services segment. Cloud based services, advanced data services, mobile backhaul, increasing video and mobile (including Cable Wi-Fi and small cells) traffic, managed services and wholesale models and larger business customers offer new revenue areas. To successfully compete in this market requires new network architectures with higher consistency in service and that support dynamic bandwidth services.

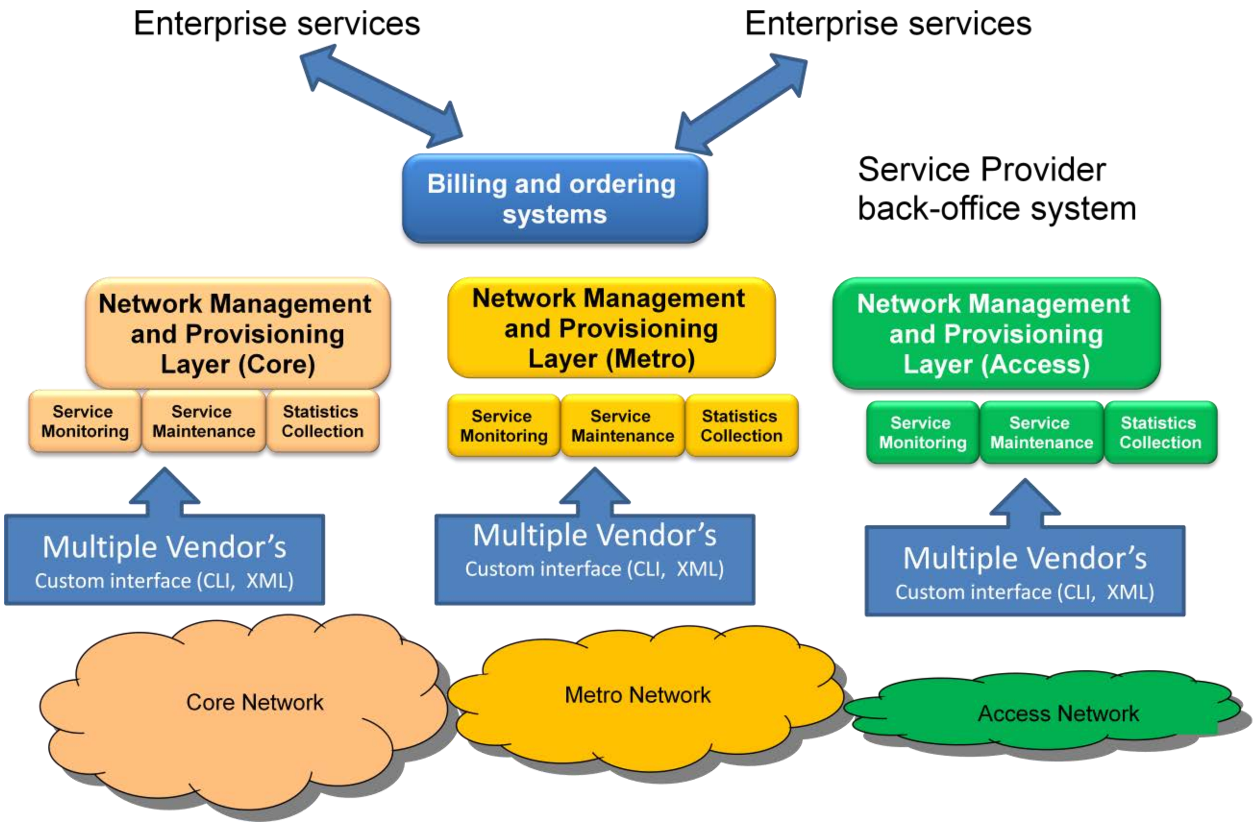

Figure 2. Segmented Network Infrastructure

While there is a common consensus on potential revenue opportunities for MSOs from video, mobile, cloud and backhaul services, a key question to be answered is "How can MSOs decrease their mean time to derive revenue while evolving to future proof systems?". To answer this question, we have to analyze current and future business services and their requirements (i.e. enterprise customer's perspective) and current and future network and system architectures with their performance and cost characteristics (i.e. service provider's perspective). The analysis should lead to evolutionary steps for Greenfield deployments and leveraging existing technology while incrementally adopting software defined and virtualized networks that may be deployed for Brownfield deployments.

In current network deployments, to provision dynamic services with end-to-end service guarantees and ensuring a SLA is not straight forward. Figure 2 shows a typical service provider network infrastructure. The current network infrastructure is divided into separate segments (Core, Metro and Access networks); each of these networks segments are controlled and managed by customized interfaces and tools. Often multiple vendor network elements and dedicated EMS/NMS systems are connected to the OSS/BSS system.

Lack of a common network control layer makes it hard to obtain a global consistent network topological view and cause the inability to instantaneously evaluate the capability of the network. Without this information it is extremely difficult to support dynamic services in a timely manner. Typically, the network has to be traffic engineered and often times there are multiple touch points (i.e. network elements) involved

– each of which need to be carefully configured and tested for inconsistencies. It involves coordination among multiple teams from these silos which is very hard to pull off in a short time period. This implies a huge cost and longer lead times that may not meet the customers' desired flexibility.

To support rolling out dynamic backhaul services to these customers requires a global network topological view and ability to evaluate the capability of the network to ensure a guaranteed service. Seamless network monitoring and performance measurement tools are required. While there are many tools and protocols, such as IEEE 802.1ag and ITU Y.1731, defined to achieve this functionality, they are non-existent in many parts of the network. This makes it extremely challenging for the operators to offer these services over today's network infrastructure.

Evolving commercial services, new business models and larger scale services impose new requirements both on network management and physical network infrastructure. For instance, network management should be flexible to support rapid service integration, dynamic bandwidth control and end-to-end SLA visibility. The underlying network infrastructure should support elastic bandwidth requirements and the scale of these new cloud based applications. There are several new standards and emerging industry groups' initiatives on providing solutions to achieve these requirements, namely Software Defined Networking (SDN) and Network Functionality Virtualization (NFV). Most of these standards originated to solve the problems in data centers and cloud infrastructure. While the network elements in data centers are quickly becoming all software defined and virtualized, service provider's networks and systems have differences that require a more evolutionary strategy.

In the next two sections, we compare the enterprise and service provider's point of views and the challenges in fulfilling the emerging requirements.

Enterprise Customer's Perspective

With the increasing adoption of Cloud for compute and storage by enterprises and businesses, application development teams have overcome many resource challenges and it has become a norm to roll out new iterations of their applications fairly frequently (reduced from weeks to days and some cases hours). Since most of these applications are based on cloud resources, they need the bandwidth and some level of QoS for their services. The IT teams in these Enterprises and SMBs are challenged to provide the necessary resources in a relatively short time period while ensuring these new applications do not have any detrimental effect on the already running critical and daily traffic.

These services along with projects that require temporary network usage, data back-ups and upgrades demand dynamic network solutions. Increasing traffic and unpredictable patterns are challenging the CIOs and IT network architects to accurately forecast the capacity requirements for the backhaul pipe. Usually a WAN service that has enough data rate capacity to satisfy the near term demands is ordered. This increases the total cost of operation for the business and often results in poor utilization of their subscribed static pipes.

These customers are actively seeking backhaul services that can be controlled dynamically instead of a static services that are either inflexible or take days to change. In addition to dynamic bandwidth changes, some enterprises are also considering more granular service options where a new business application can be deployed as a new service with end-to-end SLA visibility, monitoring and business analytics.

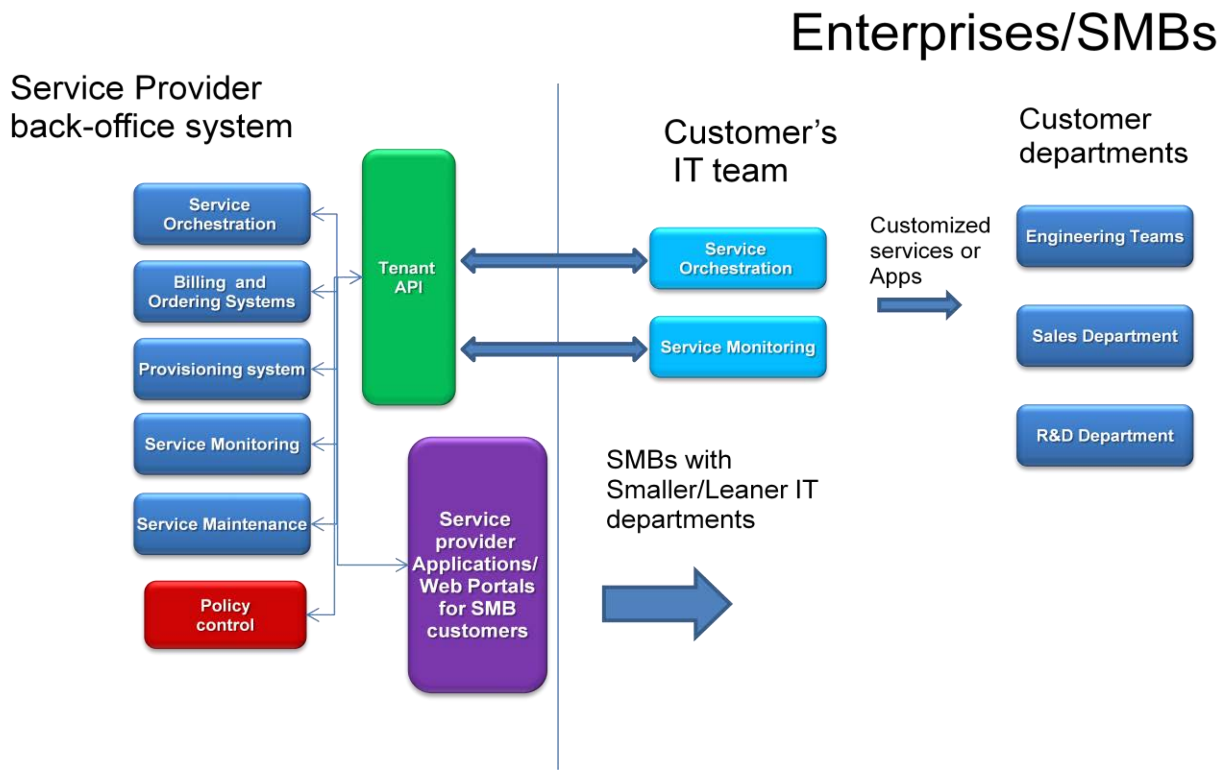

Figure 3. Exposing control to customers - Tenant API

To meet these new fast paced requirements, these customers are seeking turnkey solutions that include both technology and professional services for cloud based services; communication services including Voip, virtual PBX, video/web conferencing and other business video applications, resource and data management; managed mobile services, BYOD (bring your own device); and related security technology and services.

Figure 3 presents an ideal view of an Enterprise or SMB's IT team vision. Having the visibility, control and interface to build their own applications such as Service Orchtestration, Service Monitoring, using an open API (e.g. Tenant API as shown in Figure 3) increases the flexibility for the customers. Some SMBs and Enterprises with leaner IT teams, MSOs can offer canned applications or web portals that can be customized.

<u>Service Provider's Perspective</u>

Today, most of the backhaul connections in the access and the core network are either static and involves manual provisioning or pseudo automated provisioning. Typically the enterprise or SMB's IT team estimates the bandwidth required for each branch office and negotiate a service with the desired data capacity, SLAs and QoS requirements (e.g. 2 EVCs, CIR1 = 50Mbps, EIR1 = 15Mbps and CIR2 = 10Mbps, EIR2 = 0 etc.) from their service providers. If the customers need to upgrade the pipe capability, it requires a lot of manual intervention from both the customer and service provider. The provisioning and billing team from the service provider have to communicate and provisioning involves a lot of touch points. Unfortunately, this whole effort is in the order of days if not weeks.

Service Providers have to offer enterprise communications and managed services with the dynamic and flexible solutions as described in previous section. Other revenue opportunities such as mobile backhaul and wholesale models demand dynamic bandwidth management and end-to-end service visibility as well.

MSOs are exploring different solutions to meet high bandwidth demand by extending HFC networks (e.g. Docsis 3.1) and Carrier Ethernet and EPON networks (e.g. DPoE). In addition to physical network infrastructure, backoffice and customer management software solutions should enable flexible and dynamic service integration, control and management. Multiple MSOs should be able to serve a big enterprise with multiple locations. While supporting dynamic bandwidth requests, the new management systems should maintain high network utilization. The system provisioning, service integration and control should be simplified and independent of specific network components provided from multiple vendors as interoperability is desired.

The time scale for service introduction and bandwidth contol should be `comparable to computing and storage programmability of data centers, which may be achieved by service and network virtualization.

SOFTWARE DEFINED NETWORKING (SDN)

Software defined networking presents an interesting concept of abstracting the network infrastructure from the applications by decoupling the control and data planes. Figure 4 shows a typical SDN. Network control intelligence and state of network elements are consolidated into a centralized software server. Conceptually this software controller can be hosted on any elastic cloud infrastructure and thus can be scaled based on the demand.

SDN has emerged primarily to solve many new problems in the data center space such as provisioning and managing thousands of end points (virtual machines) that dynamically move. Data center interconnectivity requires high-capacity, low latency links that can dynamically scale based on need. The East-West communications between data centers (Inter-Data-Center) of several data intensive companies, such as Google, Facebook etc., have demanded a more configurable network that is efficient in both cost and performance. On the other hand an Intra-Data-center is characterized by highly dynamic end-points (VMs). To increase efficiencies and maintain low cost of operations, VMs are often migrated from one host to another resulting in a dynamic logical network.

OpenFlow [1] is one major attempt to radically abstract many of the networking functions. It presents a programmable interface with well defined instruction set to control and program a network. A centralized controller application can use this programming interface to configure and manipulate the network elements dynamically. Applications can use this interface to respond and adapt the underlying networking infrastructure to tackle these new challenges. A consistent and centralized logical view of the network can be obtained relatively easily.

In [3], several attributes are discussed to define SDN. Below we list the relevant attributes and show how they can help MSOs to achieve their business service objectives. Some of these attributes may be achieved by extending current technologies, leading to hybrid solutions towards the evolution of pure software defined and virtualized networks.

Logically centralized and separated control plane:

Separation of control (and management plane) from data plane enables controlling the data flow in software (without specific hardware) and control becomes logically centralized with full view of network. Today, most of the network elements in a service provider network has data and control plane in the same appliance although management plane is

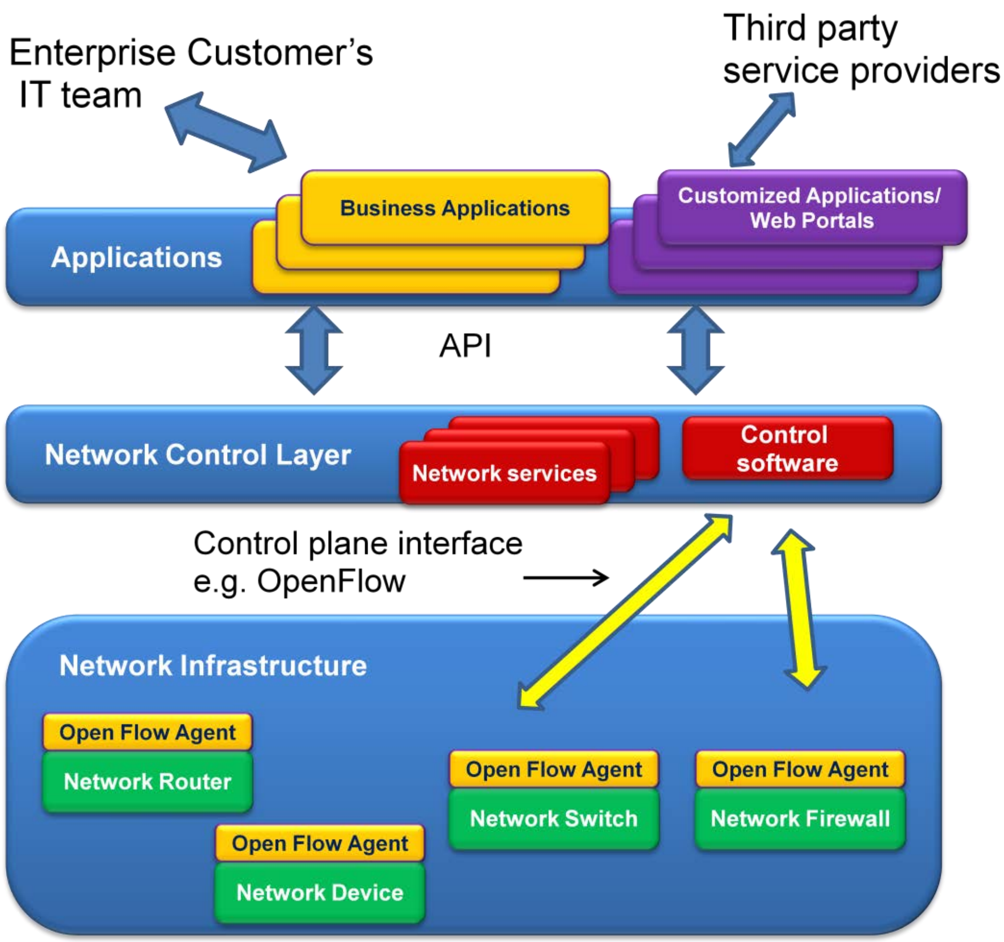

Figure 4. Typical Software Defined Network (SDN)

mostly separated. The notion of remote control is well established but appliance dependency and distributed nature are the main differences compared to logically centralized and separated control plane of SDN networks In an SP network, control of network elements may include packet, wavelengths, TDM etc. and depending on the technology the decision should be made on the control elements to be centralized. The ultimate goal for a SP is to have a modular single logical entity that is abstracted from appliances and controllable for service, subscriber and policy management.

<u>Programmability of network and service features</u>:

What makes it different than today's programmable and configurable attributes of SP's network and service elements is the time scale of dynamic change. The time scale is expected to improve over evolution of the networks. For example, a new service creation and deployment should take hours compared to weeks and be automated without dependency on a specific appliance. Programmability provides business agility by enabling enterprise customers to request dynamic changes in their services.

Network and Service Virtualization:

The feature corresponds to the abstraction of network and service attributes that can be changed dynamically through software control and dynamic programmability of the network and services. Virtual networks and services such as VPN, VLAN, VPLS, Virtual Access Points etc. help service providers to serve multiple customers and services over a single network. However, both virtualization and programmability enable automation and finer time scale compared to what virtual networks can provide today.

In addition to Open Networking Foundation (ONF), IETF, ITU's Network Funtions Virtualization (NFV [5]), Open Daylight, MEF [2] and TMFORUM [4] has ongoing efforts on providing SPs with interoperable mechanisms to extend Carrier Ethernet networks for new services. MEF introduced CE 2.0 which enables multi-CoS. (e.g. differentiating best effort Internet vs. high QoS voice for enterprise and MBH services). E-Access element provides end-to-end SLAs and management control enabling multiple cable operators to serve a big enterprise with multiple locations. E-Access also enhances wholesale Ethernet services e.g. selling to multiple wireless customers with one carrier Ethernet service. Manageability is enhanced with fault tolerance and other tools. MEF's CE4Cloud project with Dynamic Responsive Ethernet (DRE) proposal corresponds to programmability of SDN environments towards controllable software components. It addresses elastic Ethernet service attributes, including dynamic change of EVC or UNI attributes for a time period (e.g. increasing CIR of existing EVC, adding or removing UNI endpoints) and on-demand and reservation models.

ENABLING SDN and NFV

SDN and Network Functions Virtualization (NFV [5]) enable MSOs to have enhanced and dynamic service integration and management, reduce provisioning and deployment cost and time. SPs can orchestrate their network and services in an end-to-end manner from a logically centralized control with simplified provisioning and management, including dynamic service management, traffic engineering (dynamic bandwidth and load distribution), full view monitoring, fault management and business intelligenc. It provides enterprise customer and wireless service provider end-to-end SLA visibility and business intelligence analytics. Virtualization and programmability increases resource utilization.

For service providers the objective is to have both service and network virtualization. Network virtualization helps the SPs to see the network as a one entity to control and manage while service virtualization enable SP to manage, programme and automate services over the network. Virtualization does not require SDN and vice versa. Centralized control, programmability and virtualization enable the network to be movable (e.g. bandwidth allocation can be moved).

Benefits of SDN

Centralizing the intelligence and control is not a new concept for operators. This is how most of the networking has evolved over the past few decades. What SDN does differently is abstraction and consolidation of the intelligence and control portion into a centralized controller thereby simplifying and potentially eliminating the control and intelligence portion from the network elements. This offers several benefits to the operators resulting in lower OPEX and CAPEX.

With this abstraction and consolidation, hardware dependence can be eliminated and the benefits of a cloud based controller can be realized. The controller is a software module that can potentially run on a cloud environment and can be elastic and scalable as needed.

The NEs become more focused on forwarding rules and only act on the ACLs set by the controller. This can dramatically reduce the cost of these network elements.
Service providers can also achieve better efficiencies from the network

- build and maintain logical mesh networks (both in core and access) much more easily and control
- dynamic load balancing for efficiency improvements and
- network resilience can be easily achieved

A common centralized controller presents a potential for improved network provisioning, management and monitoring of the entire network

- Automated provisioing and fault correlation can bring lower operational costs.
- End to end flows, SLAs can be easily monitored.
- Statistics can be viewed internally and external access to customers can also be arranged through the API.

Time to market can be dramatically brought down yielding in quick revenue, resource optimization (virtualization, dynamic control, statistical utilization), better overall performance. (service and network planning, SLA assurance, opertational orchestration, unified billing)

A logical topology view along with the control state and intelligence at a central controller can greatly benefit the operators. Operators can obtain a global state of the network with a granular view of each flow or traffic on any link fairly easily. Adding a new service becomes a mere point and click model. Of course this assumes all the network elements support a standard protocol and present an open API for configuring. OpenFlow or similar protocol can be agreed by all vendors.

Why should MSOs care about SDN?

SDN makes a lot of sense in an environment where changes are rapid and require dynamic network reconfiguration. As can be noted from the discussions in previous sections, many IT organizations will sooner or later need to offer dynamic services to their clients. SDN will help reduce the complexity of provisioning and any manual intervention by abstracting the underlying infrastructure as new applications are deployed. This also enables the network infrastructure to be agile and highly scalable to accommodate dynamic application deployments. Another motivation to the Enterprise's IT organizations is to avoid the complex network engineering tasks (often manual planning) and leverage the toolset of SDN to simplify the QoS provisioning, ACL rule setup and dynamic monitoring of the services.

Service providers do not have to change every switch or router with Openflow (or bgp-te; PCE; IETF SDNP; ALTO) enabled virtual swicthes and neither their OSS/BSS systems have to be replaced with SDN orchestrated systems to achieve the objectives listed herein. Overlay models and intermediate control modules can be extended by examining which features are implemented in hardware and their performance and impact on service and network control. One important aspect is simplification of provisioning and operations of the system. Simplification will enable to have open APIs for automated control at both service provider and enterprise customer sites. The services and applications may be requested from the enterprise customer or application (data center/cloud) center.

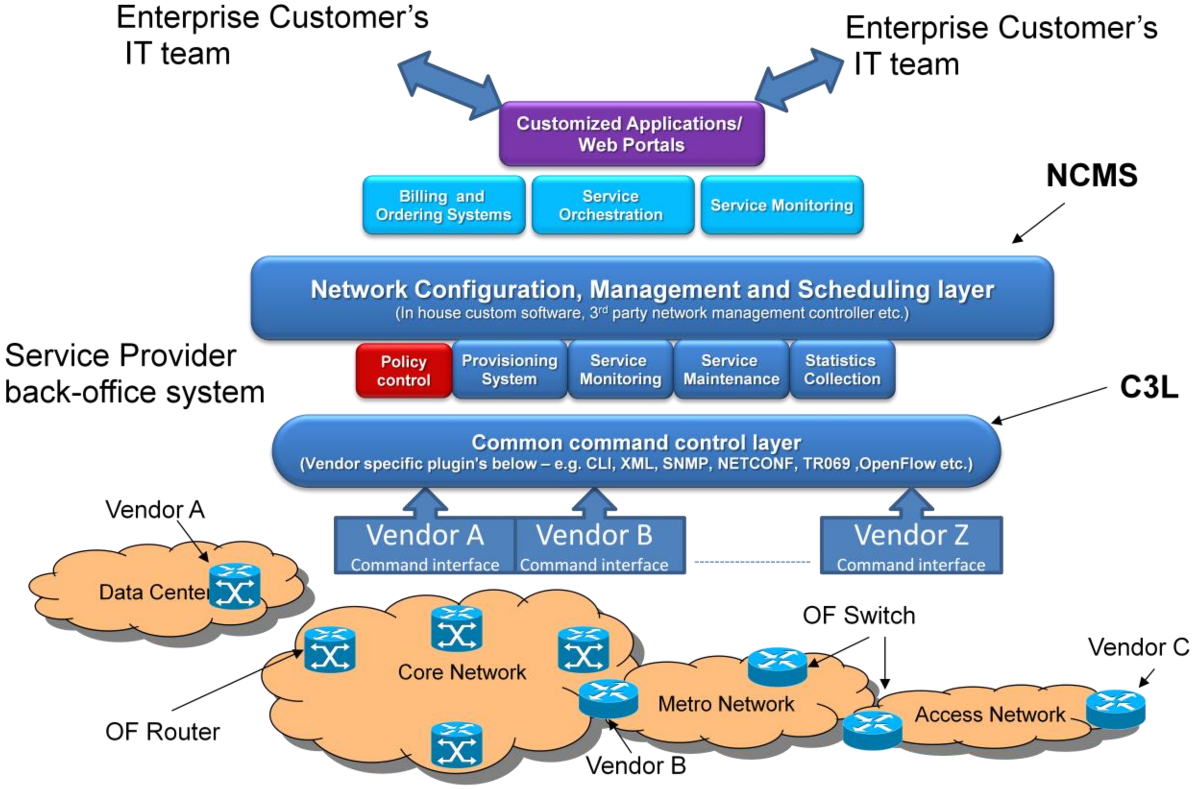

Figure 5. Provisioning Dynamic BW Services on MSOs current Infrastructure

<u>PROVISIONING DYNAMIC BANDWIDTH SERVICES</u>

This section looks at an approach of delivering dynamic services by building an overarching control layer over the existing physical network infrastructure. Figure 5 presents a high level architecture overview of the potential solution.

Accomplishing these new services such as dynamic service scheduling, providing seamless control to customers is challenging with today's network infrastructure because of lack of a common language that can be interpreted by the multi-vendor systems with various non-standard control interfaces. SNMP provides a standard configuration interface but does not completely solve the problem. Many vendors only support monitoring part of the SNMP and do not implement the configuration part of SNMP.

This can be solved by building a common command control layer (C3L) that integrates the configuration (CLI, SNMP, custom APIs) of all the vendor network elements. In this approach, vendors provide plug-in modules that integrate into the (C3L).

The abstraction in the form of NCMS vastly improves the capabilities of the operator's networks. Specifically new services can be provisioned with point and click modes even with today's network infrastructure. With this approach billing capabilities are also enhanced. As NCMS has the full view of the network, end-to-end capacity analysis may be done per customer and service. This information can be shared with the customer as available resources for dynamic service requests (e.g. using Tenant API as described below).

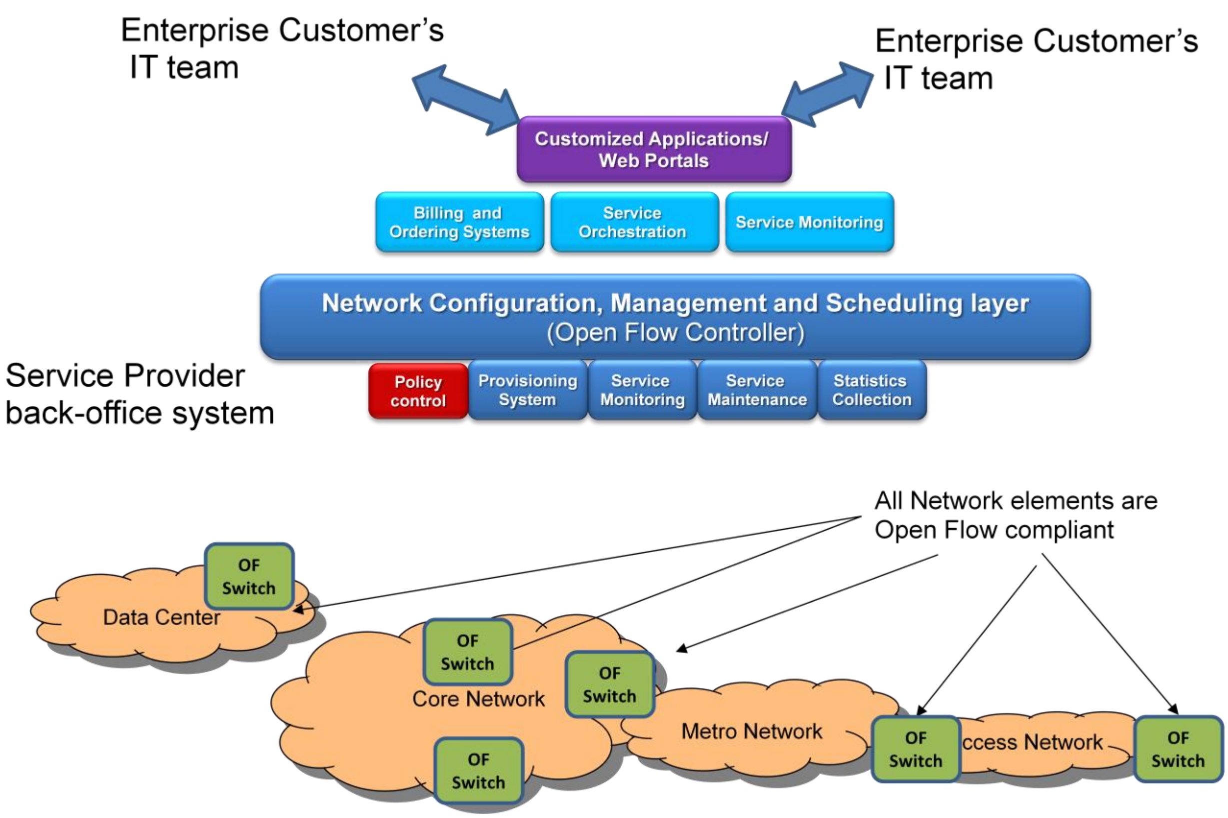

Figure 6. Pure SDN Approach

Pure SDN Approach

Ideally if all the network elements can be replaced with a OpenFlow or a similar protocol, NCMS implementation can be further simplified. OpenFlow can help by providing a standard forwarding instruction set that can configure the network elements. If all vendors adopt a common configuration control protocol then network management and control becomes a lot easier. A centralized controller requires to interface and interpret in only one language.

In addition to the advantages gained with from the approach specified in previous section, this network is much more powerful and reconfigurable. Because of the native support of an OpenFlow instruction set, all network elements can be completely controlled by the NCMS.

Furthermore, as the network elements implement OpenFlow, elimination of the complex control modules in the NEs can drive the costs of the NEs significantly down.

However, this can be very challenging and time taking effort to enable the change as it involves significant investment of resources from both the MSOs and vendor community.

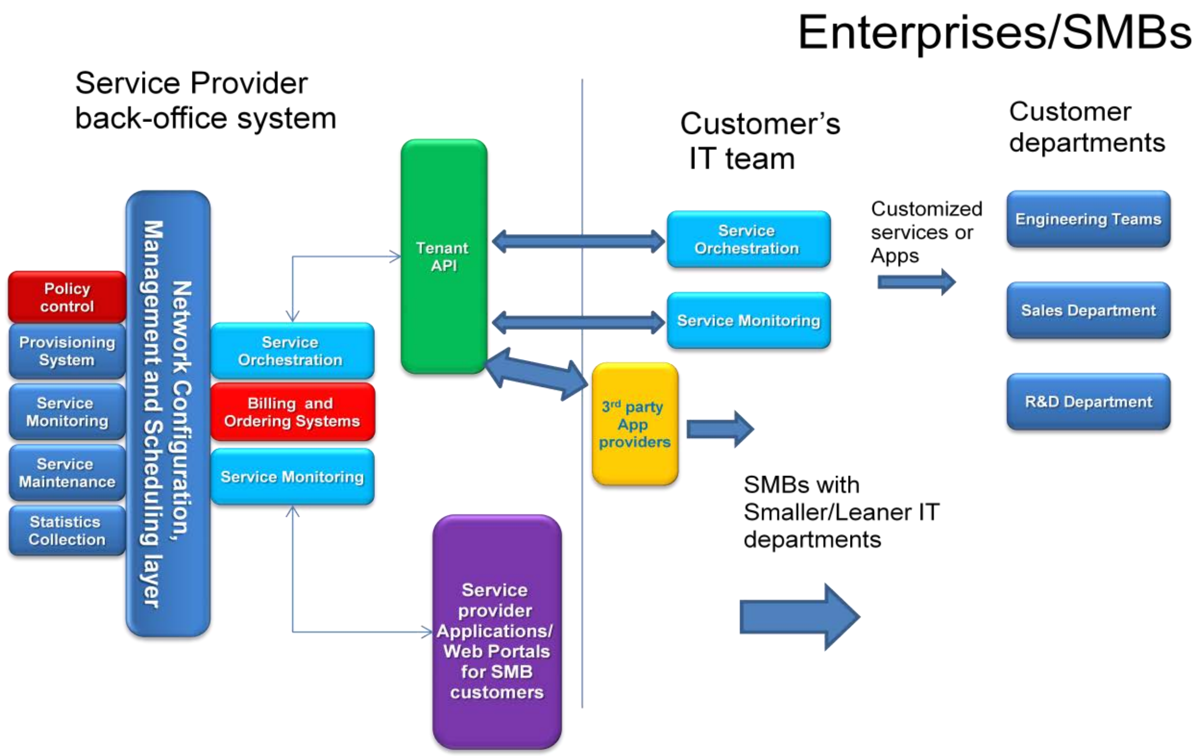

Figure 7. Tenant API: Empowering the Customers

Tenant API: Empowering the Customers

The Tenant API could allow the customers to orchestrate new services and change the data rates of existing services dynamically. Enabling this feature by the MSOs becomes easier with the NCMS abstraction layer. Background checks are made to ensure the SLAs are met and the customer can be billed appropriately for the services.

As shown in Figure 7, IT departments at larger enterprises can build their own internal applications (or buy 3rd part applications) using the Tenant API for orchestrating the services to internal teams or departments. This way they don't have to completely reveal their internal application structure or corporate details to the service providers.

In smaller or medium businesses which typically run leaner IT departments can use the applications or web portals provided by the service providers. These web portals or applications help the customers to demand customized services, control existing services and the applications will orchestrate accordingly.

With these new service orchestration, billing has to be very granular and many innovative techniques could be used to drive up the revenue from the services satisfying both the customers and service providers.

In addition to the service orchestration, service performance monitoring and maintenance features can also be provided for detailed reports on a periodic basis. These statistics can help the customers make informed decisions on future growth predictions and demand patterns. Service providers on the other hand can use this information for improving their utilization and offer services that are cost effective to the customers. This enables the customers to provision and prioritize the traffic dynamically and change over time. For example during day time certain services could be assigned high priority and during night time a different set of services are assigned with high priority.

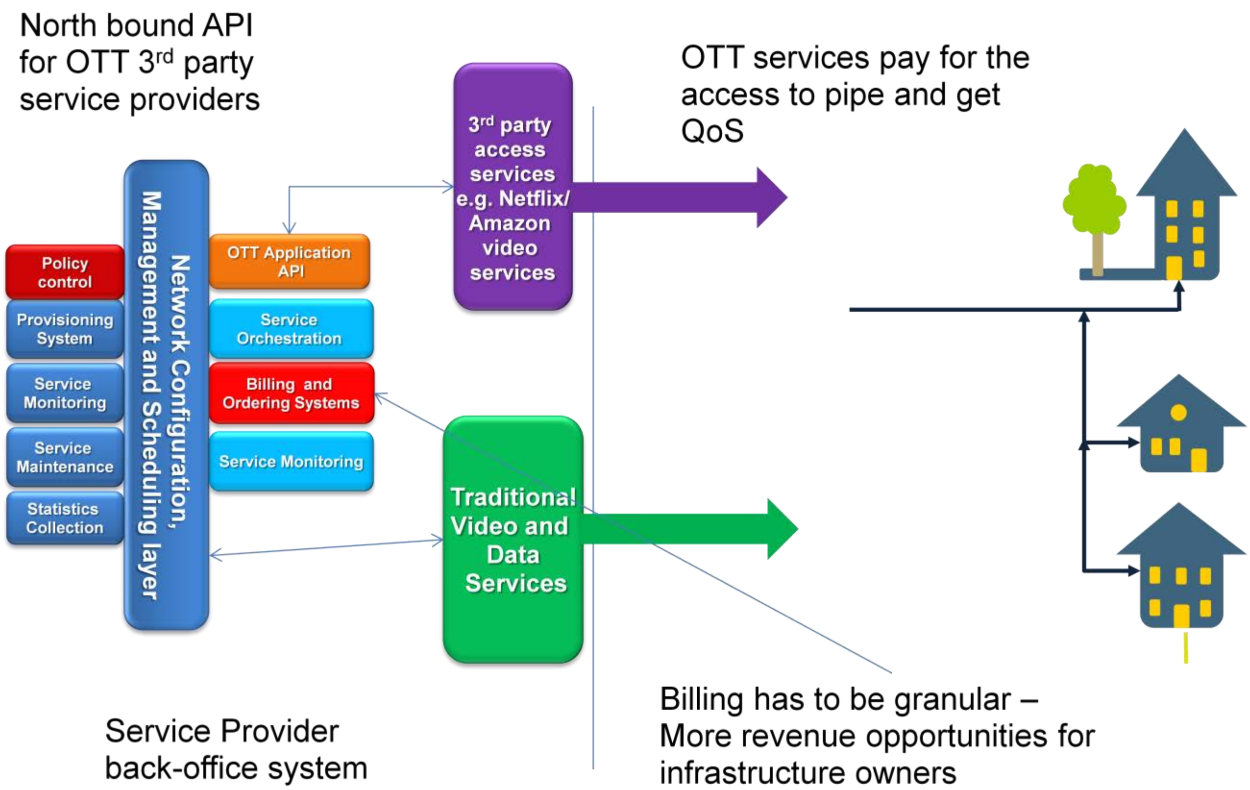

Figure 8. Virtualizing the Network Pipe

VIRTUALIZING THE ACCESS PIPE

Virtualization of anything and __ as a service is becoming has been the trend and are being embraced by many customers. Brand new business models like Network-as-a-Service (similar to Amazon's IaaS) can be explored by the MSOs. Network and Service virtualization can be defined in different ways. The existing network elements can continue to host the control plane intelligence but present Open API support to control the behavior from remote controller elements. This sort of control already exists in many devices, primarily through proprietary interfaces or by standard tunneling methods (e.g. VLANs, GRE) where the controller device sends policies to the network elements. But most of the API is either limited or closed to the vendors. Third party application developers have limited access and due to the proprietary API, development usually does not scale well. Presenting a standard API interfaces like adopting Openflow standard would present a better option for innovation.

Once a NCMS layer is built, virtualizing the current network infrastructure is straight forward. As shown in the Figure 8 an Over the Top (OTT) API can be defined that interacts with the NCMS, the service orchestration, billing and ordering systems, and service monitoring modules. The OTT API lets third party service providers to orchestrate services by requesting and reserving resources to obtain a desired QoS. The billing and ordering systems can appropriately turn on relevant counters for accounting and get paid for the services. New granular billing lets the operators to come up with innovative ways to improve the revenue opportunities.

SUMMARY

A significant number of businesses are adopting cloud infrastructure for compute and storage and application services. This has stirred the growth of new era of applications that communicate and exchange lot of data, and are deployed rapidly for testing. IT teams in these enterprises are trying to keep up with the requirements of nimble and faster response times. Soon the access network infrastructure will need to be nimble enough to match the agility of cloud data center based services and support dynamic network reconfiguration and features like multi-CoS features at a more dynamic time scale.

To address these needs, an overarching control and configuration layer is required on top of the existing physical infrastructure.

As operators make new investments into physical infrastructure, an intelligent network with dynamically reconfigurable network elements should be considered. OpenFlow appears to have a great potential with its well defined forwarding instruction set. Vendors and Service providers can start this as reference point to drive the industry to a common and intelligence control plane which can reinvigorate the network to meet the new challenges presented by the cloud based application services. This also presents a several new benefits for the service providers. First, an evolution of today's network into a virtual network can help in building new business models that monetize by selling services to third party access providers A common control and intelligence plane can help in improving the billing granularity. Services that exceed the fair use bandwidth could get fair treatment and pay for those services.

Furthermore, a network with a common network control layer will significantly reduce the time to market and will be capable of delivering new services with efficient management and streamlined operations.

REFERENCES

[1] OpenFlow [www.opennetworking.org]
[2] MEF [www.metroethernetforum.org]
[3] http://www.metanoia-inc.com/blog/
[4] TMFORUM
[5] ETSI NFV [http://portal.etsi.org/portal/server.pt/community/NFV/367]

AUTOMATING B2B ORDER MANAGEMENT FOR METRO ETHERNET SERVICES BETWEEN OPERATORS USING TM FORUM FRAMEWORX™

Brian Hedstrom
CableLabs

Abstract

Today, many MSOs rely on manual or vendor proprietary solutions for the Order Management process in delivering Metro Ethernet services. This paper proposes a next generation approach to automating the ordering process for Metro Ethernet services using the TM Forum Frameworx™ architecture, that is better suited for open, standard interfaces, common business processes, standardized information models and integrated back office applications.

This paper will develop an Order Management automated Business-to-Business (B2B) solution approach using a MEF Ethernet Private Line (EPL) plus E-Access service example use case. First, the overall TM Forum Frameworx will be introduced to provide the high level view of the Service Oriented Architecture concepts. This will be followed by the E-Access service Order Management conceptual design.

MEF EPL + E-ACCESS SERVICE USE CASE

Often an MSO (Service Provider "A") may need to access an out-of-franchise subscriber location as part of providing an end-to-end Ethernet Provide Line (EPL) service [1]. Since the subscriber for the EPL service resides in Access Provider B's footprint, this is accomplished by Access Provider B offering an E-Access service [2] to Service Provider A. Figure 1 illustrates this service topology of the EPL end-to-end Ethernet Virtual Connection (EVC) for the point-to-point link between the two subscriber locations along with the Operator Virtual Connection (OVC) that exists between the User-to-Network Interface (UNI) and External Network Network Interface (ENNI) for both the Service Provider and Access Provider. The E-Access service provided by the Access Provider is defined as OVC B between the UNI B and the ENNI AB and generally carries a Service Level Agreement (SLA). Our use case for this paper is to automate the Service Order process and interface between Service Provider A ordering the E-Access service from Access Provider B.

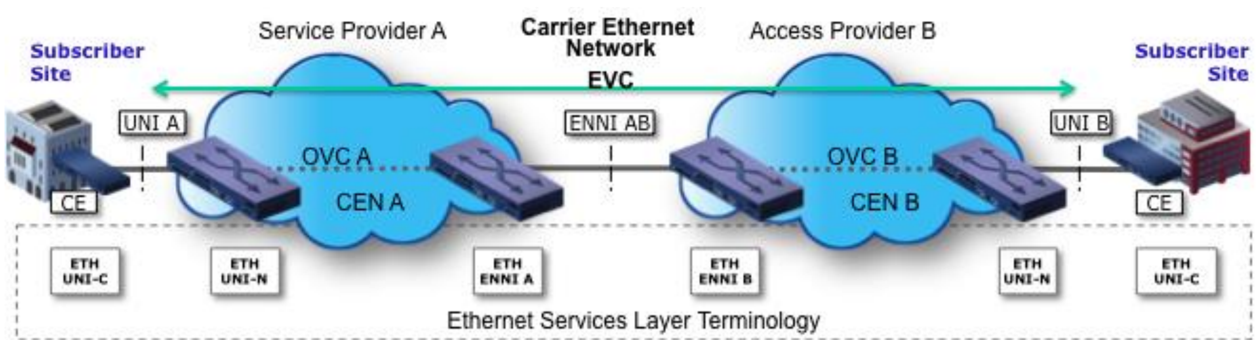

Figure 1. MEF EPL (End-to-End) + E-Access Service Order Use Case

TM FORUM FRAMEWORX

The TeleManagement Forum (TMF) has defined a Service Oriented Architecture (SOA) for Communication Service Providers called Frameworx and is diagrammatically depicted in Figure 2. Frameworx is based on defining and integrating business processes, information, applications and interfaces.

Business processes, such as Service Order Management business logic and policies, are covered in the Business Process Framework [3], which includes the enhanced Telecommunications Operations Map (eTOM). eTOM is an extensible model which defines a complex business processes blueprint for Service Providers in how they provide products to their end customers as well as how they execute their business at an enterprise level.

Information modeling, to support the flow of data between applications and interfaces to support the business processes, is covered in the Information Framework [4] and includes the Shared Information/Data model (SID). A Service Order Management Information Model, derived from the MEF Carrier Ethernet Information Model [MEF 7.2] falls into this area.

Defining a Service Order type of application for Metro Ethernet Services is covered in the Application Framework [5] and includes the TAM. The TAM defines an overall application map for Service Providers to aide in application and tool management, acquisition, inventory, etc. with a mapping to those business processes they implement in their organization.

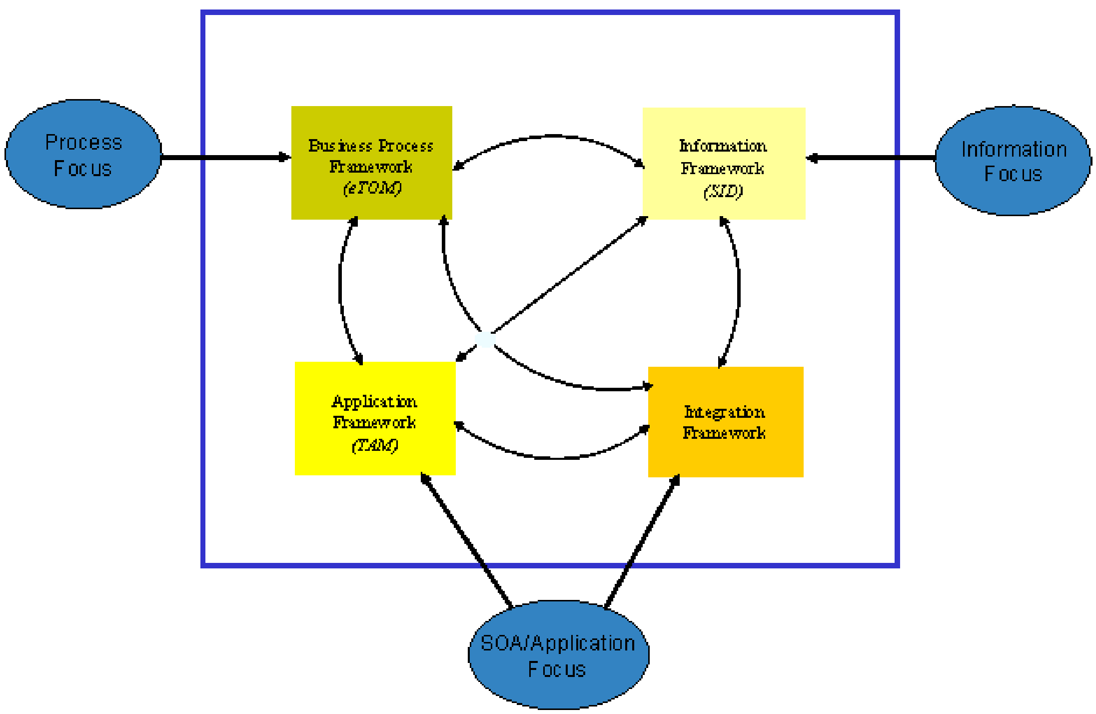

Figure 2. TM Forum Frameworx Overview

The Integration Framework binds the other Business Process, Information and Application Frameworks together. The Integration Framework includes the standard protocols, interface definitions, Application Programming Interface (API) definitions and technologies to integrate the business processes and workflows, information and data exchanges and application interfaces. For example, the Integration Framework would include a Service Order Management API definition.

BUSINESS PROCESS FRAMEWORK

The enhanced Telecommunications Operations Map (eTOM) defines a hierarchy of processes, with process decomposition from Level 0 down to Level 4. Figure 3 illustrates the high level view of the eTOM with Level 0 and Level 1 processes exposed. This view highlights three major groupings that equate to the Level 0 processes:

- Strategy, Infrastructure & Product – Covers processes that are capital expenditure centric, such as planning and lifecycle management.
- Operations – Covers processes that are operational expenditure centric, such as operational management.
- Enterprise Management – Covers processes that are specific to executing the day-to-day business activities, such as business support management. This paper will not focus on this area.
-

When addressing order management processes, both Strategy, Infrastructure & Product as well as Operations are involved. Strategy, Infrastructure & Product processes will be discussed first, followed by Operations.

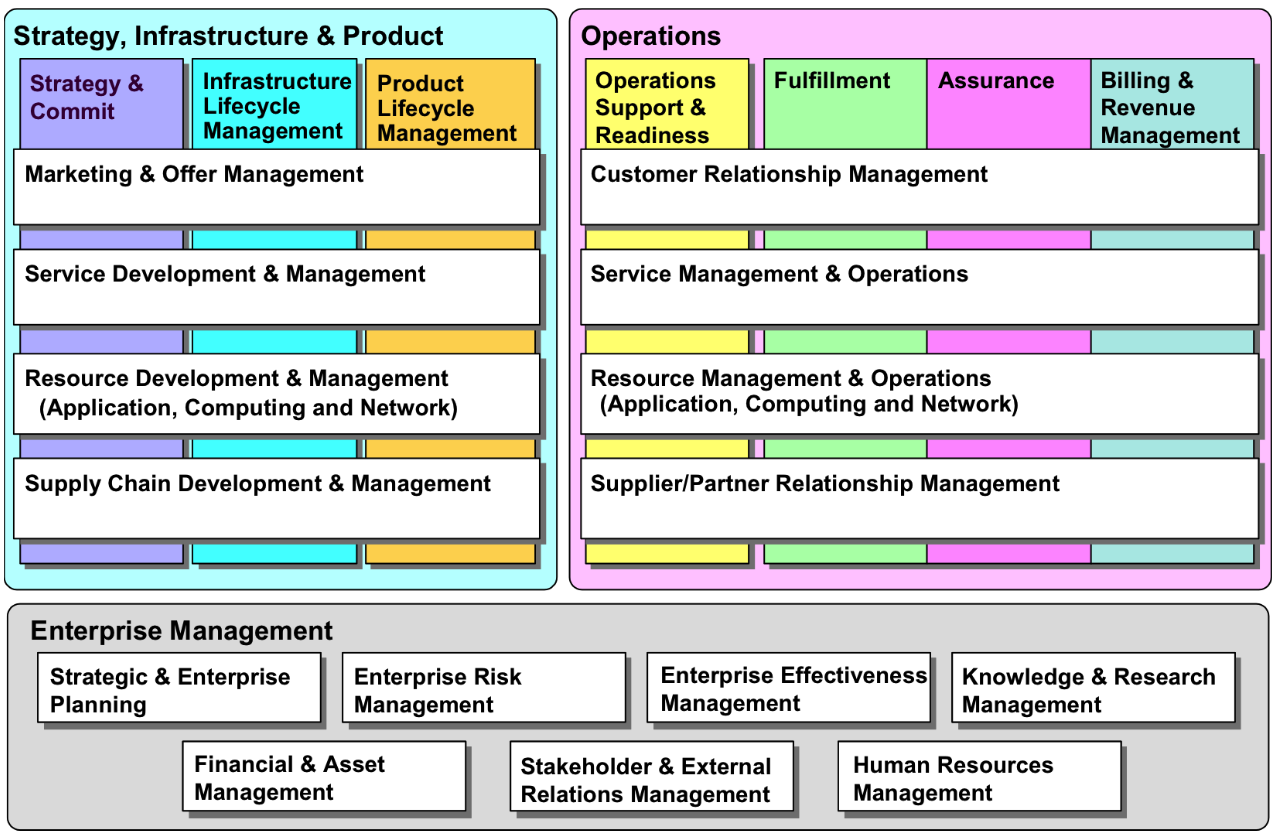

Figure 3. TMF enhanced Telecom Operations Map [3]

One important aspect of selling and buying services is to first define, plan and implement all of the services in the enterprise's product portfolio. This is often referred to as a product catalog. The product or service catalog should define the specific E-Access [2] service which the Service Provider A would like to order from Access Provider B. The service offering should have a detailed Service Level Specification (SLS) that identifies the service attributes and parameters and will help aid in developing a Service Level Agreement (SLA) with the Service Provider upon the contract negotiations. The SLS and SLA will also be inputs to turn up testing, referred to as Service Activation Testings (SAT), performed by both the Service Provider and the Access Provider. These processes are defined in eTOM [6] as identified below:

- Level 1: Product Lifecycle Management: *"This vertical end-end process grouping is responsible for the definition, planning, design and implementation of all products in the enterprise's portfolio."*
    - Level 2: Product & Offer Development & Retirement: *"Product Development & Retirement processes develop and deliver new products or services and product or service enhancements and new features, ready for implementation by the Operations processes."*
        - Level 3: Develop Detailed Product Specifications: *"The Develop Detailed Product Specifications processes develop and document the detailed product-related technical, performance and operational specifications, and customer manuals."*

With the definition of the product and service offering, via the product catalog, and the SLS completed, the majority of the other order management business processes fall into the Operations area. This area supports the customer, the network, and interaction with operations and management on a day-to-day basis. The fundamental order management business processes reside under Operations and within Customer Relationship Management (CRM). Since Service Provider A wants to purchase an E-Access service from Access Provider B, this is a CRM Order Handling business process that Access Provider B must fulfill through their own service and resource configuration and activation processes. Once they have installed, configured, activated, tested and verified the service according to the SLS, SLA and service order, they can turn the service over to Service Provider A. However, the very beginning of the process should be an automated ordering process to allow Service Provider A to order the E-Access service from Access Provider B using an application interface. These processes are covered in eTOM [6] as identified below:

- Level 1 Horizontal: Customer Relationship Management: *"This horizontal functional process grouping considers the fundamental knowledge of customers needs and includes all functionalities necessary for the acquisition, enhancement and retention of a relationship with a customer."*
- Level 1 Vertical: Fulfillment: *"The Fulfillment process grouping is responsible for providing customers with their requested products in a timely and correct manner."*
    - Level 2: Order Handling: *"Order Handling processes are responsible for accepting and issuing orders."*
        - Level 3: Determine Customer Order Feasibility: *"Check the availability and/or the feasibility of providing and supporting standard and customized product offerings where specified to a customer."*
        - Level 3: Authorize Credit: *"Assess a customer's credit worthiness in support of*

managing customer risk and company exposure to bad debt."

    - Level 3: Track & Manage Customer Order Handling: *"Ensure customer provisioning activities are assigned, managed and tracked efficiently to meet the agreed committed availability date."*
    - Level 3: Complete Customer Order: *"Manage customer information and interactions after customer contracts or associated service orders have been finalized and during the order completion phase."*
    - Level 3: Issue Customer Orders: *"Issue correct and complete customer orders."*
    - Level 3: Report Customer Order Handling: *"Monitor the status of customer orders, provide notifications of any changes and provide management reports."*
    - Level 3: Close Customer Order: *"Close a customer order when the customer provisioning activities have been completed."*
- Level 2: Service Configuration and Activation: *"Allocation, implementation, configuration, activation and testing of specific services to meet customer requirements."*
    - Level 3: Allocate Specific Service Parameters to Services: *"Where the Allocate Specific Service Parameters to Services processes are requested by a service order issued in response to a confirmed customer order, these processes are responsible for allocating the specific service parameters required to satisfy the initiating service order."*
    - Level 3: Issue Service Orders: *"Issue correct and complete service orders."*
    - Level 3: Close Service Order: *"Close a service order when the service provisioning activities have been completed."*

The eTOM defines many Level 4 processes for Order Management. However, this paper will not get into that level of process detail. Based on this level of order process detail residing in eTOM, no extensions are necessary for our use case.

INFORMATION FRAMEWORK

The core of the Information Framework, or Shared Information/Data model (SID), is information modeling. Information is a fundamental thread when designing interfaces, protocols, APIs, applications, business process workflows, etc. As our use case is specific to a Metro Ethernet service, the MEF has developed their own industry specific Information Model. An Information Model is the design of the objects and parameters in any system, including operations and actions, as well as how different objects are associated with one another (e.g., their assocations). Information Models are commonly defined using UML class diagrams, which the TM Forum and MEF have adopted.

When looking at the Order Management information, we need to address both the TM Forum SID from an industry agnostic perspective and at the MEF Carrier Ethernet Management Information Model [7] for the industry specific perspective. The TM Forum SID provides us with the generic information about order management and associates the information with the eTOM business process definitions. The Frameworx has been designed to be generic as possible so that it can be used across any industry (government, telecommunications, energy, etc.) However, a generic interface will not provide us with a MEF E-Access service order solution, which is were we need to integrate the industry specific information from the MEF Carrier Ethernet Management Information Model. This essentially is applying the Information Framework for Service and Resource Management for Ethernet Services.

The SID defines a Business Entity, an Aggregate Business Entity (ABE) and a Domain. Domains are the horizontal collection of ABEs associated with specific management areas that correlate directly to the eTOM horizontal business processes (e.g., Service, Resource, etc.). An ABE is a well-defined set of information and operations that characterize a highly cohesive, loosely coupled set of business entities (e.g., Service Configuration) [8]. A Business Entity is simply something of interest to the business (e.g., customer, customer order, customer account). Think of a Business Entity as a managed object or class in a class diagram. Figure 4 illustrates the SID with the Domains and Level 1 ABEs exposed.

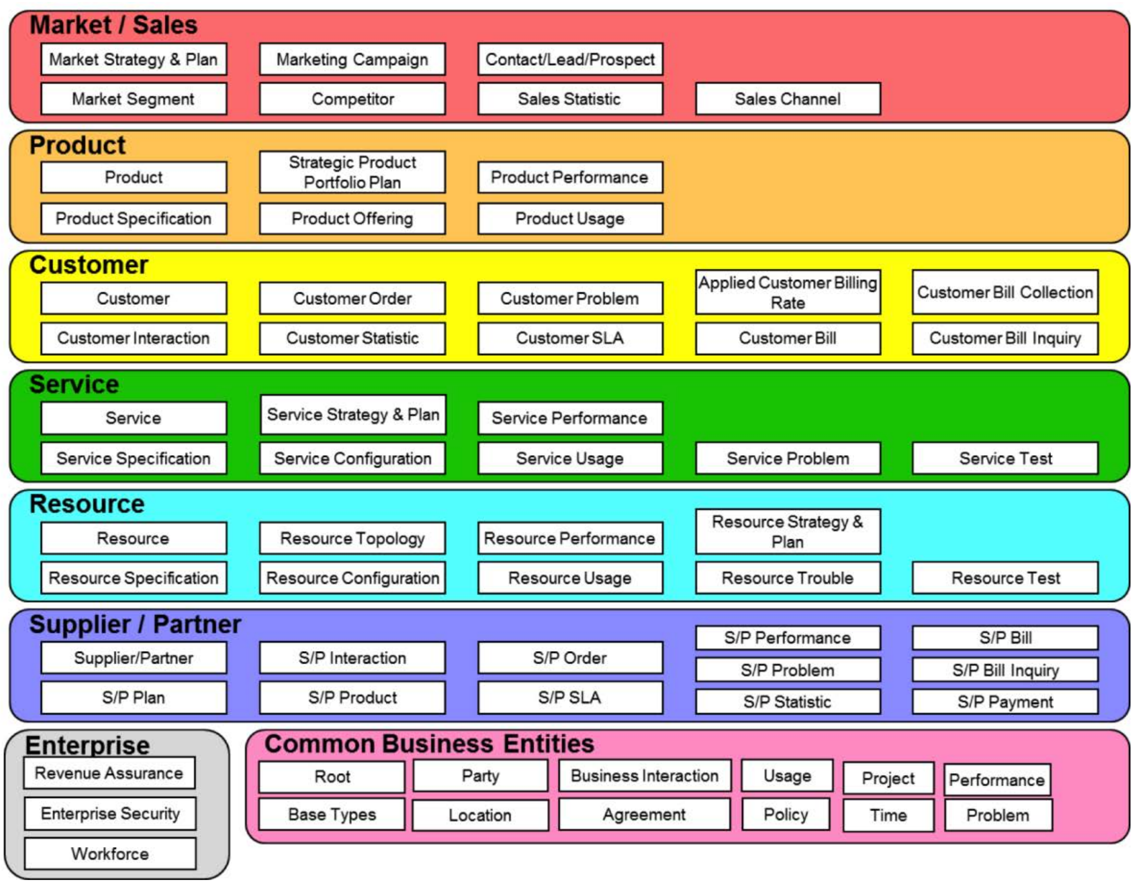

Figure 4. Information Model (SID) including Domains and Level 1 ABEs [8]

The SID Level 1 and Level 2 ABEs and Business Entities [8] of interest for Order Management include:

- Product
- Product Specification
- Product Offering
    - Product Catalog ABE (Level 2)
- Customer Interaction
- Customer Order
    - CustomerOrder/CustomerOrderItem Business Entity
    - ProductOrder/ProductOrderItem Business Entity
- Service
    - Service Order ABE (Level 2)
        - ServiceOrder/ServiceOrderItem Business Entity
- Service Specification
    - Customer Facing Service Specification ABE (Level 2)
- Service Configuration

These are further decomposed into lower level ABEs and Business Entities, however this paper will not get into that level of detail.

The MEF E-Access service attributes and parameters are modeled in ITU-T Q.840.1 [9] and MEF 7.2 [7]. This pair of documents combined defines the complete set of Carrier Ethernet Management Information Model as well as defining use cases to support the UML class diagrams used to model the information. In order to develop an automated interface for an E-Access service order, we need to create a service order view into the Carrier Ethernet Management Information Model. This identifies which objects and attributes are necessary to support the order entry process. For example, the Carrier Ethernet Management Information Model includes the managed objects necessary for performing Service Operations, Administration and Maintenance (Service OAM) for assisting in network fault and performance monitoring. However, such managed objects are not required for the service order entry process. A few of the OVC service attributes include:

- End Point List/Map
- Max MTU Size
- CE-VLAN ID Preservation
- CoS Identifier
- CoS Name
- Ingress Bandwidth Profile per OVC End Point at a UNI

There are many more E-Access service attributes, however this paper will not get into that level of detail.

APPLICATION FRAMEWORK

As discussed in the previous sections, the Business Process Framework (eTOM) provides a view of business processes while the Information Framework (SID) provides a view of information and data for Service Providers; the Application Framework (TAM) provides a view of applications. These applications are typically software-based, procurable, deployable, OSS and BSS level applications or services. One goal of the Application Framework is to define a common language set for information exchange based on eTOM when referencing operation and business support systems. Other goals of the Application Framework include standardizing application requirements and automation enablement through componentization of support systems. The purpose is to bridge the gap between the eTOM and SID and applications. The top level Application Framework is shown in Figure 5.

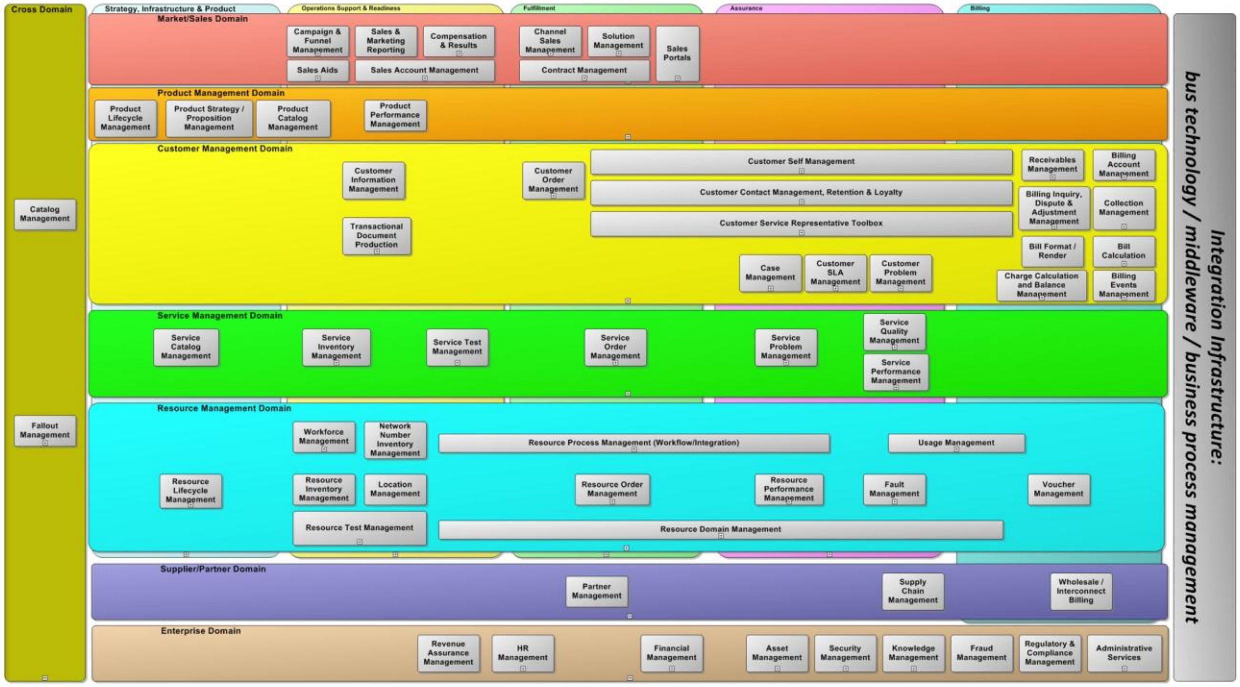

Figure 5. Application Framework (TAM) [5]

The application areas specific to the Order Management use case should align with the eTOM business processes and SID Level 1 and 2 ABEs and Business Entities already identified and described in the previous sections. These application areas include [5]:

- Product Management Domain
    - Product Lifecycle Management
    - Product Catalog Management
- Customer Management Domain
    - Customer Order Management
- Service Management Domain
    - Service Order Management

The above set of application areas provide a detailed but generalized set of application requirements to address the MEF E-Access Order use case. Refer to [5] for the detailed description and requirements for each area. Each application area can be extended with additional requirements that are specialized for the E-Access MEF service. For example, the Order Management might be extended with a requirement that requires the Access Provider to supply a copy of the OVC test report showing their Service Activation Test results for their E-Access service offering.

INTEGRATION FRAMEWORK

Finally, the Integration Framework is the cohesion of all three of the other frameworks. It defines a set of standards that supports interoperability between applications defined in the TAM while based on the requirements from the eTOM. The interfaces are defined in terms of standardized information and data models from the SID. The result is a reusable Service Oriented Architecture approach to addressing the Service Providers enterprise's needs. It has long been recognized that by standardizing interfaces to OSS and BSS applications, Service Providers realize an operational efficiency through reduced integration costs, improved interoperability, reliability, and agility. The MEF E-Access use case requires an inter-company business-to-business (B2B) OSS Order Management client interface. Java EE™ provides three open technology platforms (through the OSS though Java Initiative OSS/J) for this interface scenario. The TMF defines a corresponding integration profile [10] for each interface technology. These are listed below:

- Remote Method Invocation (RMI) over IIOP
    - Uses the OSS/J Java Profile (JVT)
- XML over Java Messaging Service (JMS)
    - Uses the OSS/J XML/JMS Profile
- Web Services (SOAP over HTTP)
    - Uses the OSS/J Web-services Profile

The TM Forum, via the OSS/J, has developed an Order Management API (JSR-264) [11] to enable creation of a standardized interface for applications. The API specifies the requirements for Order Management system interfaces based on the eTOM and SID as described in the previous sections. The API is also extensible to allow the addition of the E-Access service attributes defined in the Carrier Ethernet Management Information Model [7]. In addition, operations are defined to create, abort, modify, remove, suspend and cancel orders and order activities [12]. Figure 6 illustrates the E-Access use case in detail.

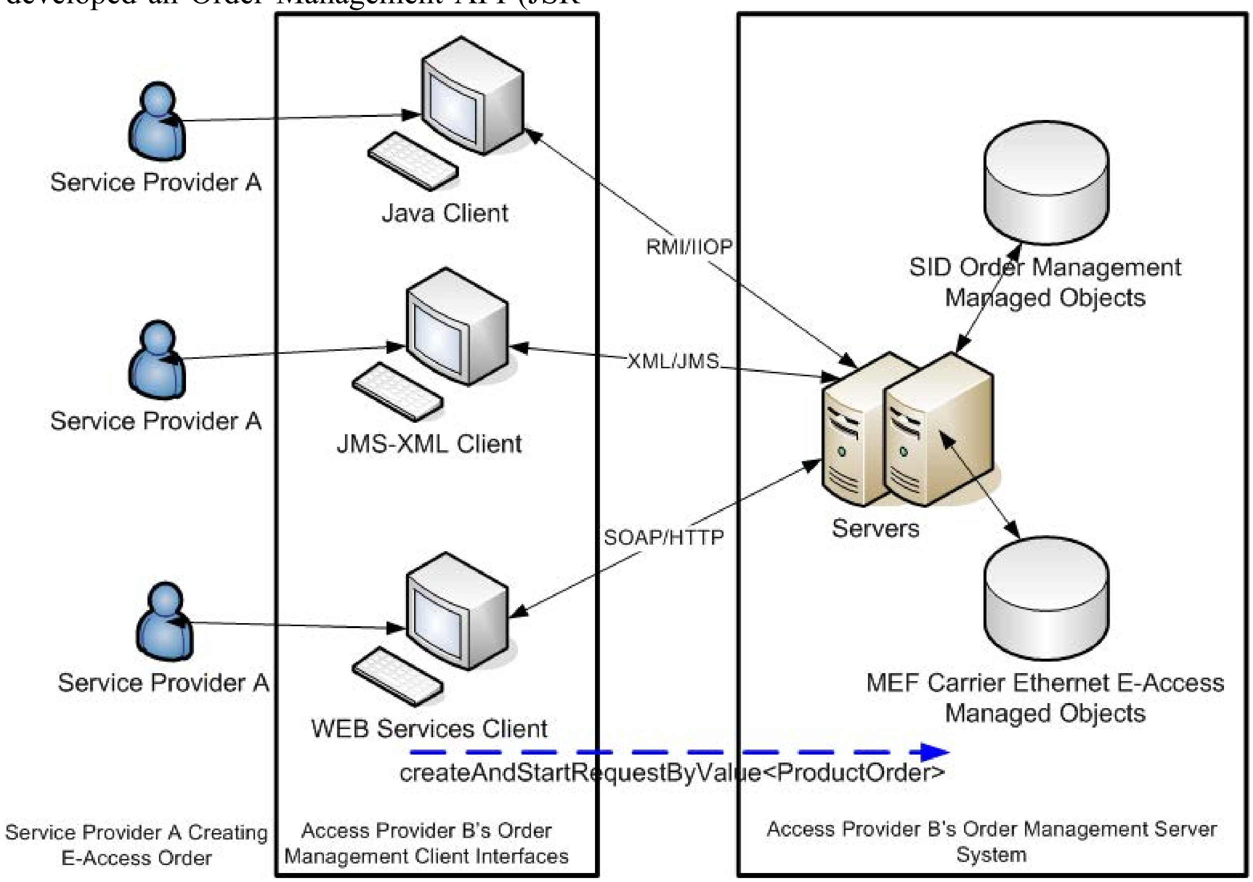

Figure 6. Programmable Order Interface for E-Access Use Case

Access Provider B's Order Management system is composed of a client-server application which, in the simplest form, accesses the SID managed objects, such as the Product Offering from the Product Catalog. The E-Access service attributes are also extensions to the SID managed objects, as these describe the service specific details beyond the generalized SID model. The Order Management server can expose the three client side interfaces: RMI/IIOP, XML/JMS, and SOAP/HTTP. The Web

Services (SOAP) interface has become the dominate interface in the industry and is the focus of this section.

Service Provider B accesses the Order Management application via a client side interface, such as the Web-services client to *create* an order for the E-Access service, based on the Product Offer derived from the Access Provider's Product Catalog. Once the Service Provider fills in the required attributes and submits the order, the hashed blue arrow in Figure 6 illustrates the SOAP message `createAndStartRequestByValue` sent to the server containing the `<Product Order>` details [13]. When the server receives this message it begins the necessary Service Order, including the Service Specification details, and Service Configuration processes discussed earlier in the paper. This may invoke other applications, interfaces, etc. internal to the Access Provider's back office. Appendix I shows an example XML message sent in the SOAP message containing some of the SID managed objects.

SUMMARY

This paper has presented a Service Oriented Architecture approach to designing an automated order management interface for ordering Metro Ethernet services between operators by leveraging the TM Forum Frameworx. This was accomplished by instantiating each of the following core areas within Frameworx:

- Business Process Framework

The MEF Service Order Business Processes/Business Logic/Business Policies were identified.

- Information Framework

The Order Management SID entities were identified. The MEF Service Attributes and Parameters for the E-Access Product and Service Offering were identified as extensions to the SID.

- Application Framework

The B2B Operation Support System/Business Support System (OSS/BSS) applications for Order Management were identified, based on the business processes and information of the eTOM and SID.

- Integration Framework

Finally, The OSS/J Order Management Application Programming Interface (API) was leveraged to design and implement a standard web-services SOAP interface for creation of a Order Management client-server system for use in allowing the Service Provider to programmatically create and submit an order to the Access Provider's system.

BIBLIOGRAPHY

[1] MEF 6.1 Technical Specification, Ethernet Services Definitions – Phase 2, April 2008.

[2] MEF 33 Technical Specification, Ethernet Access Services Definition, January 2012

[3] Business Process Framework (eTOM), Concepts and Principles, GB921 Version 9.2, TM Forum, April 2011.

[4] Information Framework (SID), Concept and Principles, GB922, Version 7.3, January 2012.

[5] Application Framework (TAM), GB929, Version 4.7, TM Forum, April 2012.

[6] Business Process Framework (eTOM), Addendum D: Process Decompositions and Descriptions, GB921 Addendum D, Version 12.2, TM Forum, April 2012.

[7] MEF 7.2 Technical Specification, Carrier Ethernet Management Information Model, July 2012.

[8] Information Framework (SID), User's Guide, GB922, Addendum 1U, Version 8.6, TM Forum, September 2011.

[9] ITU-T Recommendation Q.840.1, Requirements and Analysis for NMS-EMS Management Interface of Ethernet over Transport and Metro Ethernet Network, 2007.

[10] OSS/J Design Guidelines, Part 1 – Basic Concepts and General Design Patterns, Release 2.0, TM Forum, April 2008.

[11] OSS/J Order Management API, JSR-264 Overview, Release 1.0, TMF892, January 2010.

[12] OSS/J Order Management API, JSR-264 User's Guide, Release 1.0, TMF893, January 2010.

[13] OSSJ-OrderManagement-v1-0.xsd, XML Schema File provided in OSS/J Order Management API, JSR-264, Release 1.0

APPENDIX I: EXAMPLE ORDER CREATION XML MESSAGE


```xml
<?xml version="1.0" encoding="UTF-8"?>
<createAndStartRequestByValueRequest xmlns="http://ossj.org/xml/OrderManagement/v1-0"
xmlns:cbebi-v1-5="http://ossj.org/xml/Common-CBEBi/v1-5" xmlns:cbecore-v1-
5="http://ossj.org/xml/Common-CBECore/v1-5" xmlns:cbedatatypes-v1-
5="http://ossj.org/xml/Common-CBEDatatypes/v1-5" xmlns:cbelocation-v1-
5="http://ossj.org/xml/Common-CBELocation/v1-5" xmlns:cbeparty-v1-
5="http://ossj.org/xml/Common-CBEParty/v1-5" xmlns:cbeproduct-v1-
5="http://ossj.org/xml/Common-CBEProduct/v1-5" xmlns:cbeproductoffering-v1-
5="http://ossj.org/xml/Common-CBEProductOffering/v1-5" xmlns:cberesource-v1-
5="http://ossj.org/xml/Common-CBEResource/v1-5" xmlns:cbeservice-v1-
5="http://ossj.org/xml/Common-CBEService/v1-5" xmlns:co-v1-
5="http://ossj.org/xml/Common/v1-5" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
instance" xsi:schemaLocation="http://ossj.org/xml/OrderManagement/v1-
0:\PROGRA~1\altova\xmlspy2006\schemas\ossj\OSSJ-OrderManagement-v1-0.xsd">
    <requestValue xsi:type="ProductOrderValue">
        <co-v1-5:key>
            <co-v1-5:type>ProductOrder</co-v1-5:type>
            <co-v1-5:primaryKey>order-123456</co-v1-5:primaryKey>
        </co-v1-5:key>
        <cbebi-v1-5:interactionDate>2001-12-17T09:30:47.0Z</cbebi-v1-5:interactionDate>
        <cbebi-v1-5:interactionDateComplete>2001-12-17T09:30:47.0Z</cbebi-v1-
5:interactionDateComplete>
        <cbebi-v1-5:interactionStatus_BusinessInteraction>unknown</cbebi-v1-
5:interactionStatus_BusinessInteraction>
        <cbebi-v1-5:description>String</cbebi-v1-5:description>
        <!-- the locations involved in this order -->
        <cbebi-v1-5:involvedLocations>
            <cbebi-v1-5:item>
                <cbebi-v1-5:place>
                    <co-v1-5:lastUpdateVersionNumber>2147483647</co-v1-
5:lastUpdateVersionNumber>
                    <co-v1-5:key>
                        <co-v1-5:type>String</co-v1-5:type>
                        <co-v1-5:primaryKey/>
                    </co-v1-5:key>
                </cbebi-v1-5:place>
            </cbebi-v1-5:item>
        </cbebi-v1-5:involvedLocations>
        <!-- the (party) roles involved in this order -->
        <cbebi-v1-5:involvedRoles>
            <cbebi-v1-5:item>
                <co-v1-5:key>
                    <co-v1-5:type>String</co-v1-5:type>
                    <co-v1-5:primaryKey/>
                </co-v1-5:key>
                <cbecore-v1-5:describedBy>
                    <cbecore-v1-5:item>
                        <cbecore-v1-5:value/>
                        <cbecore-v1-5:characteristic>String</cbecore-v1-
5:characteristic>
                    </cbecore-v1-5:item>
                </cbecore-v1-5:describedBy>
                <cbebi-v1-5:interactionRole>provider or subscriber</cbebi-v1-
5:interactionRole>
            </cbebi-v1-5:item>
        </cbebi-v1-5:involvedRoles>
        <om-v1-0:priority_Request>0</om-v1-0:priority_Request>
```

```
        <om-v1-0:expectedCompletionDate>2001-12-17T09:30:47.0Z</om-v1-
0:expectedCompletionDate>
        <om-v1-0:validFor>
            <cbedatatypes-v1-5:startDateTime>2001-12-17T09:30:47.0Z</cbedatatypes-v1-
5:startDateTime>
            <cbedatatypes-v1-5:endDateTime>2001-12-17T09:30:47.0Z</cbedatatypes-v1-
5:endDateTime>
        </om-v1-0:validFor>
        <om-v1-0:requestedCompletionDate>2001-12-17T09:30:47.0Z</om-v1-
0:requestedCompletionDate>
        <om-v1-0:purchaseOrder>String</om-v1-0:purchaseOrder>
        <!-- each productorder item is effectively a line item in the order -->
        <om-v1-0:productOrderItems>
            <om-v1-0:item>
                <co-v1-5:key>
                    <co-v1-5:type>String</co-v1-5:type>
                    <co-v1-5:primaryKey/>
                </co-v1-5:key>
                <!-- the actual action being requested - e.g. "create" the product or
"cancel" the product or "modify" the product  -->
                <cbebi-v1-5:action>Create</cbebi-v1-5:action>
                <cbebi-v1-5:quantity>
                    <cbedatatypes-v1-5:amount>1</cbedatatypes-v1-5:amount>
                    <cbedatatypes-v1-5:units>String</cbedatatypes-v1-5:units>
                </cbebi-v1-5:quantity>
                <!-- the specific places for this productorder item - will usually
refer to an involvedlocation specified on the "parent" order level -->
                <cbebi-v1-5:places>
                    <cbelocation-v1-5:item>
                        <co-v1-5:key>
                            <co-v1-5:type>String</co-v1-5:type>
                            <co-v1-5:primaryKey/>
                        </co-v1-5:key>
                </cbebi-v1-5:places>
                <!-- the specific roles for this productorder item - will usually refer
to an involvedroles specified on the "parent" order level -->
                <cbebi-v1-5:involvedRoles>
                    <cbebi-v1-5:item>
                        <co-v1-5:key>
                            <co-v1-5:applicationContext>
                                <co-v1-5:type>PartyRole</co-v1-5:type>
                                <co-v1-5:primaryKey>subscriber-123</co-v1-5:primaryKey>
                        </co-v1-5:key>
                        <cbebi-v1-5:interactionRole>Subscriber</cbebi-v1-
5:interactionRole>
                    </cbebi-v1-5:item>
                    <cbebi-v1-5:item>
                        <co-v1-5:key>
                            <co-v1-5:applicationContext>
                                <co-v1-5:type>PartyRole</co-v1-5:type>
                                <co-v1-5:primaryKey>provider-123</co-v1-5:primaryKey>
                        </co-v1-5:key>
                        <cbebi-v1-5:interactionRole>Provider</cbebi-v1-
5:interactionRole>
                    </cbebi-v1-5:item>
                </cbebi-v1-5:involvedRoles>
                <!-- the details of the product (offerring) being ordered -->
                <om-v1-0:productOffering>
                    <co-v1-5:key>
                        <co-v1-5:applicationContext>
                            <co-v1-5:type>ProductOffering</co-v1-5:type>
```

```
                            <co-v1-5:primaryKey>productoffering-123456</co-v1-
5:primaryKey>
                    </co-v1-5:key>
                    <!-- these name value pairs are characteristics and are used for
flexible non-schema defined attributes -->
                    <cbecore-v1-5:describedBy>
                        <cbecore-v1-5:item>
                            <cbecore-v1-5:value>value1<cbecore-v1-5:value/>
                                <cbecore-v1-5:characteristic>attribute1</cbecore-v1-
5:characteristic>
                        </cbecore-v1-5:item>
                        <cbecore-v1-5:item>
                            <cbecore-v1-5:value>value2<cbecore-v1-5:value/>
                                <cbecore-v1-5:characteristic>attribute2</cbecore-v1-
5:characteristic>
                        </cbecore-v1-5:item>
                    </cbecore-v1-5:describedBy>
                    <cbeproductoffering-v1-5:description>Ethernet Access
service</cbeproductoffering-v1-5:description>
                    <cbeproductoffering-v1-5:name>E-Access EPL</cbeproductoffering-v1-
5:name>
                    <cbeproductoffering-v1-
5:state_ProductOffering>active</cbeproductoffering-v1-5:state_ProductOffering>
                    <!-- here would go the MEF-Access product service specific
attributes -->
                    <!-- <mef-v1-0:attributeXYZ>value1<mef-v1-0:attributeXYZ/> -->
                </om-v1-0:productOffering>
            </om-v1-0:item>
        </om-v1-0:productOrderItems>
    </requestValue>
</createAndStartRequestByValueRequest>
```

BIG DATA: CAPITALIZING ON UNTAPPED KNOWLEDGE

Marthin De Beer and Kip Compton
Cisco Systems, Inc.

Abstract

Immense stores of information essential for effective, efficient and profitable decisions are distributed throughout today's cable plant. Most of it is untapped, under-used or useless because there are no collection, selection and visualization tools to help assess its value. Yet every day, cable operators could leverage Big Data to solve a myriad of problems—from distribution plant issues and capacity planning to improving the customer experience and deriving higher revenues from effective advertising.

This paper:

- *Outlines Big Data,*
- *Provides examples of Big Data archtectures,*
- *Describes sources of Big Data in existing systems,*
- *And explores Big Data's relevance for cable operators.*

BIG DATA: DEFINITIONS AND LANDSCAPE

For any large company, innovation holds the key to growth and future success. In the near future, a typical enterprise will rely on new and different knowledge gleaned from *Big Data* to innovate, compete and grow in creative ways—and to do so quickly. Big Data is data that is too large, too distributed and comes in too many disparate formats to process and understand using traditional methods.

The 2013 Cisco Connected World Technology (CCWT) Report shows that more than two-thirds of IT managers see Big Data as a strategic priority for their company in 2013 and beyond.[1]

Avalanche of Big Data

Consider the sheer volume of data. Servers in data centers have traditionally used business software to access information stored in databases in disk drives on storage frames. These databases usually hold between several gigabytes (GB) and several terabytes (TB) of data. Big Data, however, requires databases that can handle a petabyte or more at a time.

Today, traditional databases often cannot manage the the amount of data to be explored. By 2012, Walmart was already handling more than 1 million customer transactions every hour, imported into databases estimated to contain more than 2560 TB or 2.5 petabytes (PB) of data.[2] By 2008, Google was already processing 20 PB a day.[3] In fact, the entire volume of business data worldwide across all companies is estimated to double every 1.2 years,[4] with an increasing number of enterprise data mining efforts exceeding traditional database capacity.

Works well with new kinds of data

The variety of data available today in the Internet of Things is also challenging traditional mining methods. For example, smartphones, sensors, video cameras, smart meters, GPS and social media are generating enormous amounts of data that can yield valuable insights. The 175 million daily tweets in the world[5], for example, might contain important findings about consumer perception of a company and its products. Three in four companies say their Big Data strategy will include analysis of data from these sources.[6]

Critical for time-sensitive analysis

The speed of capturing critical information is becoming more important as well. Akamai, for example, analyzes 75 million events per day to better understand targeted advertisements.[7] Financial institutions attempting to catch credit card fraud need to identify suspicious transactions within minutes of them being made.

A changing landscape

In contrast to databases that sum values to produce results, Big Data sets may constantly change. The data can be anywhere and can be of variable quality or usefulness. The sets will frequently contain *unstructured data*, for example, images, email, videos, and documents in one set. This is a big contrast to the *structured data* sets commonly found in traditional databases, that are defined by *schema* (rigid blueprints for how a database will be structured).

Big Data is a powerful trending tool, that relies on intelligent people to constantly refine it: forming a hypothesis, building a model, validating it then making a new hypothesis. It requires specially trained "data scientists" to interpret visualizations, key in interactive queries and develop algorithms that all uncover meaningful findings.

Implications for leaders

Big Data is also changing how and when decisions are made and needed. The ability to take smaller risks and get near real time feedback allow for the rapid evolution of decision making. Called A/B testing, this technique tries multiple options in rapid succession or in parallel, to gauge user or system response.

By allowing multiple assumptions to be tried, assessed and compared simultaneously, companies that have Big Data capabilities can evolve strategic decisions continuously.

Leaders will need to trust findings from Big Data—and their Big Data staff experts--in order to remain competitive in the years to come.

These are all reasons why Big Data requires a robust and secure architecture that is very different from traditional data warehouses.

BIG DATA ARCHTECTURE

Big Data systems have four key components that can make these huge data capture projects easier and more productive: collection, storage, analysis and visualization.

Collection: streaming vs. batch

The type and function of the device generating the data determines how it is collected. Key decisions to make first are:

- determining the type of collection (streaming or batch),
- the rate of collection,
- and impact on the infrastructure used for data collection.

Let's examine streaming versus batch collection. Batch collection and update are the same as in traditional systems that collect and cache data then update the database. Streaming databases are relatively recent innovations, where the database is constantly changing as updates are continually coming in. If a project requires near real-time information and decisions, it will also require streamed collection of data.

Storage remains a critical part

Since storage architecture decisions have the greatest impact on hardware costs, it is important to understand how the dataset will be queried, the frequency of the queries and how fast the results are needed. High speed and frequency require large memory architectures. A wide spectrum of queries may require a tiered architecture, with memory and storage size as the key variants.

Analysis: The heart of Big Data

Analysis and analytic tools are what uniquely define Big Data implementations. Understanding the problem and aligning the right algorithms to extract an optimum solution set make all the difference. The right analytics determine whether your processing produces a response of 64 million entries or 64. Big Data is akin to mining diamonds; the object is to extract the very few precious gems and discard the many tailings that surround them.

There are multiple types of analytics engines both in market and undergoing research. The most fascinating are learning engines. Similar to voice recognition systems, these analytics tools use past searches to assist in future analysis. Their ability to chain together filters, or perform mathematical functions on entire sets of results—using exclusive OR (XOR), for example—offers startling insights on vast pools of data.

Visualization pinpoints the expected, uncovers the unexpected

Even after analytic tools extract Big Data it may not be in a useable form. Big Data queries rarely seek "one" answer, and often the "right" dataset to a query might still be multiple megabytes in size.

After visualization systems identify trends in Big Data, or exceptions to a trend—then leadership can comprehend, analyze and act on the findings.

Cataloguing, representing and humanizing the vast stores of information that Big Data systems are capable of ingesting has become a hotbed of current research. A rapidly changing set of personalizable tools (like 3D graphics, color nesting and active elements) will allow the analyst to grasp what a few years ago would have been uncomprehensible.

Multiple types and vendors of databases

When Google mapped the Internet in the early 2000's, the company was one of a group that pioneered software tools to make searches more efficient. It is even more vital today, as the internet is now about 700 million sites and a trillion pages. These innovations were the beginning of the Big Data movement.

The open source community mirrored these early innovations over the following decade. The Apache Hadoop software project is a major software resource in Big Data. Hadoop contains components such as file systems, job schedulers and MapReduce for parallel processing of data sets. Associated with Hadoop are other open source projects such as Cassandra (fault tolerant database), Chukwa (distributed data management), HBASE (large table management), among many others.

Big Data principles and cable

Now let's consider the volume and types of data generation in the cable industry. The cable plant has numerous sources of data: Head end components, fiber nodes, switched digital video (SDV) switches, set top boxes, video on demand servers, and content delivery network servers, among others.

SDV switches are a good example of the amount of data these sources can generate, Let's assume:

- Each SDV service group served 250 subscribers.
- Each data collection is 100KB.

- And data collection was done each hour.

Then each SDV service group would generate 2.4MB of data per day.

Considered at a macro level, for each million subscribers, SDV switches alone would generate 9600Mb or 9.4GB of data *each day*. If each STB were to generate 100KB of RF, network, user and content data per hour (a very conservative estimate), each million subscribers would generate an astonishing 2400GB, or 2.34TB of data per day.

So each year, a cable operator with 25 million subscribers could generate 87600GB (85.55TB) of SDV data and 876TB of STB data.

The call center is another valuable data source for the cable operator. Call center knowledge combined with network information can predict potential issues with specific components, problem areas with lines or poles or modify maintenance practices to reduce network downtime.

BIG DATA BENEFITS FOR CABLE OPERATORS

When structured databases are enough

When field technicians are dispatched to a customer premise, often they go in unaware of the environment and conditions they will need to diagnose. As such, they will frequently fix the symptoms of the problem, but will not cure the root cause of the customer complaint. This often leads to follow-on customer calls, truck rolls when neighbor properties exhibit the same problem and often lower customer satisfaction of service.

What could happen if a field technician were to understand the RF environment instead of simply installing an amplifier every time a channel came in fuzzy for a customer?

The data exists, but most often is not collected unless a customer raises an issue. In the scenario above, if the field technician were to be told that in the last month RF SNR was down 40 percent across an entire group of homes, the technician would start at the fiber node rather than in an individual home, solving numerous reported or soon-to-be-reported problem tickets.

Let’s take this example one step farther. What if the technician was dispatched before any customer complaint occurred because trending data triggered an alarm that after house A, SNR has exceeded thresholds, impacting house B through K. Truck rolls could roll out in a scheduled and deliberate rather than reactive manner, call center volume would decrease and customer satisfaction would remain unaffected.

The example above outlines a problem that a standard structured database and traditional analytic engine could solve.

Big Data analytics in cable: The field

Let’s examine a case where a structured database will not work. A technician is dispatched to fix an intermittent error: “Every afternoon lately the video quality is really bad. Gets better in the late evening though.”

Ideally the technician would want to know all the RF characteristics for the house and surrounding properties for the last week. The technician notes that the weather had been heating up recently, so the natural language question would be: *Show me all the RF issues in the area when the temperatures exceeded 90 degrees?*

This would allow the technician to isolate the cause—perhaps a wire stretch/sag issue, or an overheating component issue that could have broader future service implications. This natural language question on data in various structures across large storage domains was a primary reason Big Data was created.

Big Data analytics in cable: Viewership

Operational efficiencies are not limited to the field arena. One of the problems cable operators grapple with is channel grouping in the numerous properties they service. One logical question to ask the data might be: *Show me the top 30 channels watched between 6pm and 9pm weeknights in the following zip code.* This could differentiate what channels should be switched versus broadcasted.

Further refinements could include *Show me the channels least watched during midnight and 6am weekends*, or depending on duration of data stored, *Show me which channels move from the top 30 watched to the bottom 10 watched during summer prime time hours.* Once the data is in the system, the only limitation is the degree of the analytics interface programmability.

Big Data is useful when various sources of information are blended together. This is key when it comes to understanding viewing habits. What if data could blend weather and channel habits? *Who watches channel A when it rains outside?* Or *What channels are popular on snow days?* What if viewership over holidays is important? *What channels are least watched on Memorial Day?*

Big Data analytics in cable: Advertising

Big Data can help generate additional advertising revenue. Data sources throughout a cable plant can identify user location, time and content selection—and quickly correlate further user choices of linear, streamed or Video on Demand. The 85.55TB of SDV data from the earlier example could easily be mined to provide detailed near real-time or trended knowledge of subscriber, channel and viewing habits.

Let's look at this information in a different way via channel affinity. Big Data shines when asking multi-dimensional questions. It may be valuable for cable operators to know, for example: *What channels do users who watch Channel A more than 80 percent of the time Monday to Friday between 6 and 7 watch Saturday at 1pm?* Or *When sports program Z airs, what Video On Demand requests are most popular?* A clear map of subscribers, their content affinities and preferred mode of consumption gives cable operators a powerful tool in advertising negotiations.

Actually monetizing advertising requires other tools beyond simply using Big Data findings as negotiating leverage. Depending on the size of the subscriber base with the same content affinity, cross selling of advertising is a clear-cut method to increase advertising impact and ad revenue.

As the subscriber set becomes more targeted, however, cable operators will need tools such as dynamic ad insertion. Big Data then allows ongoing feedback on whether ads were viewed or skipped through, if they caused channel change or if they were reviewed. This is valuable information for content producers and advertisers alike.

Big Data analytics in cable: Capital investments

Understanding when, where and how to upgrade, replace or discard existing investments is a common challenge in the service provider business. Big Data can provide valuable insights in this process. The basic premise is to augment standard metrics, views and methods with real-time information on the performance, reliability and impact of existing investments.

Current evaluation methods can help with important decisions like:

- Modularizing blanket upgrades
- Proactively changing components that impact new services about to be launched.
- Modifying or repositioning interconnects, peering and transit points based upon subscriber or session feedback.

Big Data could influence all capital decisions from vehicular use and reliability to call-center equipment selection.

CASE STUDY: USING BIG DATA TO IMPROVE CDN PERFORMANCE

Many factors are driving service providers to move to an IP CDN architecture (diagram below). Key among them is the emerging trend of multi-screen viewing.

Tablets, mini-tablets and smartphones have enabled mobile video viewing, social video interaction and provided a personalized always-on consumption platform.

Today's consumers want to watch and interact with what they want, when they want it and wherever they are. The resulting variable quality hot-spot and mobile network connections generate numerous bursts of video demand and non-linear viewing habits as well as affecting the CDN network.

Another challenge cable operators face is consumer quality requirements. Customers do not hold "over the top" content providers to the same quality standards as cable operators—yet want the connectivity, flexibility and device support that OTTs provide.

An IP CDN brings flexibility and elasticity to the CDN to deal with these new challenges. In comparison to traditional static CDN technologies, however, the dynamic environment required may affect streamed quality.

So the question is, *How to ensure quality video from the IP CDN?*

Diagram 1 below outlines the components and layers of an IP CDN. The upper tier consists of a centralized ingest and storage layer. The middle tier is a massively scalable caching layer, with the lower tier providing highly optimized edge streaming capabilities.

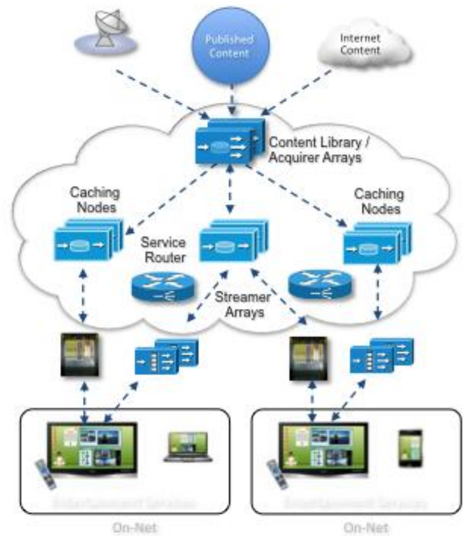

Diagram 1. IP Content Delivery Network

Small peaks cause big problems

User experience issues are either systemic issues or transient. Systemic issues are much easier to fix as basic diagnosis clearly identifies the solution to the problem. Transient issues are more complicated and require higher fidelity of information and analysis to flush out the root problems.

With an IP CDN, transient durations are typically in the two-to-three second range and can come from a number of sources. An important question to answer is, Is it the result of a sudden burst of user requests? If not, how did the issues occur?

Some possible causes are:

- Catalog server issues – did a customer request a link that was unavailable?
- Network issues – client ingress or CDN interconnect issues?
- Storage issues – sufficient storage or cache miss?
- Server issues – origin, streamer or ingest server problems?
- Or something else, such as physical unplug, power surge, or network attacks?

The Conviva Q1 2013 Viewer Experience Report states in an analysis of 22.6 billion streams roughly 60 percent of all streams experienced quality issues.[8] The three main user experience control points are:

- Buffering
- Video start time
- Grainy, low resolution picture quality due to low bit rates.

The Conviva report also notes that a 400 percent increase in viewing abandonment occurred if video start time exceeded 2 seconds.[9] Viewers who received higher bit rates watched 25 percent longer.[10]

In order to provide the best viewing experience possible one needs to know where the issues occured, when and for how long.

Collecting data sets

As can be seen in diagram 1 the IP CDN has numerous levels and sources of information. A data set to provide solutions for IP CDN could also include client types (browser, Apps…), streaming protocol used (Apple, HLS, Microsoft Smooth Streaming, Adobe HDS, Progressive Download and more…), client ISP/ geography, download size, requests per content and average percentage of content viewed.

As with the examples in the earlier sections, a typical IP CDN generates approximately 100GB daily of log data for each 100Gbps of CDN. This data only represents input from inside the CDN. If client data would be recorded as well, the storage requirements would be orders of magnitude more.

Network and content analysis

Once the appropriate data sets are collected, the analysis can focus on network and content perspectives.

Let's first look at the network view. Queries like these can glean key information: "*For time x to y, show me any interfaces with 90 percent or greater utilization.*" Or *"For this content stream, show me the utilization rate of all of the streamers."* Textual or visual representations of this data can highlight transient fault isolation.

The broader service provider network also has a major impact on the CDN. Consider, for example, rerouting, peering or congestion issues. Combining broader network information and with the CDN network data will allow multi-level questions: *If CDN egress A is less than 70 percent utilized, tell me what peering point C utilization was at time Z.*

Sometimes the issue is off-net and the network issue occurs outside of the CDN. This is when client participation in gathering statistics and usage becomes key. Client data can determine if certain ingress streaming nodes were overwhelmed or load balancers need tuning to level the requests, or if client requests from certain networks need throttling.

By combining CDN network, operator network, off-net and client information cable operators can create a detailed network scorecard that provides a comprehensive view of the transport conditions.

There are two methods to tune the CDN: Manually or through an event engine. In the manual method, IT staff would perform the analysis and tuning. An event engine is capable of both near real-time analysis of the CDN data and of tuning the CDN parameters within preset limits.

Content usage information can also provide numerous insights about the CDN. Understanding the affinities between client types, request time and location or client—or viewing time and/or percentage and client geographical concentration—provides vital information. After it is partitioned into scorecards, this information yields valuable trend analysis. Of these, the key metric of interest is: *What was the bit rate for these streams over the entire session duration?* This will indicate whether the CDN is offering the content at the quality levels customers are demanding.

The CDN could leverage the consumption patterns to preposition popular content, free-up cache space or proactively prioritize client requests, in anticipation of sudden load.

The methods described above leverage the data gleaned mainly from the inside of the CDN. What about the data that could be mined from the client perspective? A cable operator could map multi-screened subscriber patterns so that they could act proactively—from the client request moment all the way through the consumption session.

Here is an example. If a network location was prone to network congestion, session drops or other issues caused by ongoing use, the CDN could increase the client video buffer, and have the client request secondary network access or lower the streaming rate. These and other proactive actions could greatly enhance the customer experience. Bringing together client, network and service information and participation is the only way to make this possible.

Big Data helps monetize demand

Cable operators can gain a fresh strategic view of their CDNs by combining the network and content scorecards. They could proactively use trending analysis to pre-position or stage hardware assets and distribute content assets in anticipation of peak usage. They can tune continual iterations of their analytics to predict content adoption patterns. All of this information is key to both operational groups, and advertising teams who intend to best monetize demand.

CONCLUSION

The Big Data movement has evolved from research in the early 1990's, to startups in the early 2000's to today's maelstrom. Where does it go from here?

Big Data is still an island unto itself in most instances. Operationalizing Big Data systems, that is, linking and leveraging them into the operations of the business, will add greater operational efficiency and bottom line growth. Extending the knowledge and capabilities of Big Data into applications, both mobile and cloud based, will allow cable operators to offer new and dynamic services, and to quickly adjust them.

Big Data will have a dramatic impact on operational velocity, as it allows for near real-time feedback. Organizations taking full advantage of this information will become increasingly dynamic in micro experimentation. In the past, enterprises planned any large service, large application, broad release with great diligence, far in advance, to offset the immense risk. Within a

Big Data architecture , they will be able to target and launch multiple concurrent services in a fast feedback environment that mitigates risk and shortens time to revenue.

Analytic algorithm advancements will enhance effective rationalization of larger and more diverse data sets in realistic timeframes. As existing operational problem sets are mapped into algorithms, the numbers and fidelity of the analysis will increase rapidly.

The infrastructure to support Big Data is rapidly evolving—from large memory compute platforms to optimized storage strategies that allow for ever-increasing data sets. As distributed analysis techniques evolve, better analysis partition and data location in distributed data centers will optimize existing network links.

Big Data allows for a broader understanding of the subscriber base and appreciation of the unique user environments that will enhance user experience—and differentiate service providers.

The entire Big Data architecture is also evolving rapidly. In the future, learning systems will programmatically determine what data the system "thinks" it needs to collect, where to distribute it and modify the algorithms that will parse it.

As we come to understand Big Data problems better, system controllers or eventing engines will automate and execute the visualization and analysis steps, leaving the more complex and esoteric issues to humans.

What's next?

Multi-screen consumption is only the first hint of how dramatically mobile technology will change service providers. As cellular systems advance, users are rapidly breaking down perceived barriers of what and where services could be offered. What mini-van today doesn't have one or more fixed screens in it?

Consumption and user session transfer between portable, home and vehicular all offer exciting opportunities. Consider how this might transform the average family roadtrip. Typically seen as "four-wheeled torture," it could instead could become an interactive journey, where locations, attractions, vehicles, content and people all come together for a truly enjoyable experience. All we need is the data to stitch it all together.

Service providers can best use Big Data in their deployed applications to discover new and different dimensions in device and user environments. Relevant data collection will give operators a deeper understanding of error conditions, content preferences and advertising differentiation.

While the greater complications and logistical issues of extracting Big Data from these sources is huge—so is the opportunity to glean new knowledge and actionable metrics. Big Data contains priceless information that will improve operational efficiencies, enable new and personalizable services, enhance the customer experience and unearth revenue opertunities for cable operators over many years to come.

REFERENCES

[1] *Cisco Connected World Technology Report* (CCWTR), IT Manager Report 2012. Retrieved April 23, 2013 from cisco.com http://www.cisco.com/en/US/solutions/ns341/ns525/ns537/ns705/ns1120/Top-10-Survey-Results-CCWTR-Big-Data.pdf

[2] SAS. *Big Data Meets Big Data Analytics*. Retrieved April 23, 2013 fom wikibon.org

http://www.sas.com/resources/whitepaper/wp_46345.pdf

[3] Erich Schonfeld, "Google processing 20,000 terabytes a day and growing," *TechCrunch*, January 9, 2008. Retrieved April 23, 2013. http://techcrunch.com/2008/01/09/google-processing-20000-terabytes-a-day-and-growing/

[4] "eBay Study: How to Build Trust and Improve the Shopping Experience," *knowit,* W.P. Carey School of Business, Arizona State University. May 8, 2012. Retrieved April 23, 2012. http://knowwpcarey.com/article.cfm?cid=25&aid=1171

[5] Keith Dawson and D. Daniel Ziv, "A Conversation on the Role of Big Data in Marketing and Customer Service," *CRM*, April 25, 2012. Retrieved April 23, 2013. http://www.mediapost.com/publications/article/173109/a-conversation-on-the-role-of-big-data-in-marketin.html#axzz2R2jydjUA

[6] (CCWTR), IT Manager Report 2012. http://www.cisco.com/en/US/solutions/ns341/ns525/ns537/ns705/ns1120/Top-10-Survey-Results-CCWTR-Big-Data.pdf

[7] Jeff Kelly, "Taming Big Data (A Big Data Infographic), *Wikibon Blog,* May 21, 2012. Retrieved April 23, 2013. http://wikibon.org/blog/taming-big-data/ .

[8] *Conviva 2013 Viewer Experience Report*, February 13, 2013, pg. 3. Retrieved April 25, 2013. http://www.conviva.com/vxr2/

[9] Ibid.

[10] Ibid.

Breathing New Lifespan into HFC: Tools, Techniques, and Optimizations

Dr. Robert L. Howald
ARRIS

Abstract

The persistence of compound annual growth rates for IP services coupled with the rapid expansion of HD has the industry considering upgrade paths for it most valuable asset – the HFC network. HFC has an unparalleled history of flexibility, consistently adding new services over decades of graceful evolution. Nonetheless, as powerful as the network is, it remains far from optimally using its capacity resources. Given the dynamics of growth and relentless nature of compounding traffic, the inefficiency of precious HFC resources must be dealt with. New initiatives in DOCSIS and IEEE 802.3bn (EPoC) are determined to do just that.

Fortunately, modern digital communications technology has snuggled up quite closely to theoretical capacity bounds. Capacity starts at the bottom – the bottom of the OSI Stack, that is. Analogous to what Internet Protocol (IP) is for Layer 3 in the OSI stack, Layer 1 in the non-baseband world is "standardizing" around go-to Physical Layers tools such as multi-carrier technology and advanced forward error correction (FEC). These tools push performance against Shannon bounds so tightly that there is little room for further optimization. Cable is poised to take advantage of these breakthroughs to squeeze the most from HFC.

In this paper we connect the dots between network performance and network capacity. We describe application of "standard" Layer 1 tools, and how novel implementations of these tools optimize and breathe new lifespan into the HFC network. A fundamental philosophical break from the past is making use of the fact that network capability varies across users. Optimization approaches not only update the QAM technology itself, but selectively apply modulation "profiles" to make the best use of available capacity. We reveal analysis of QAM potential and profile optimization against key system variables: standard and distributed access architectures, premise architecture, network geography, component performance, and spectrum allocation. We also address the why, how, and implications of kicking the 30-year old downstream/upstream spectrum asymmetry habit, and enumerate additional results for evolved band plans in both directions. Lastly, we introduce a recently developed comprehensive, updated HFC channel model. These models go beyond prior lower-limit only references, capturing new limits, but more importantly describing "typical" scenarios and performance. A representation for "typical" channels enables determination of the ability of a network to "scale up" in capacity.

From this discussion, the audience will gain a renewed appreciation of the long runway for HFC, and a thorough understanding of the core elements of the technology upgrade. They will understand key trade-offs among technology, architecture, and spectrum relative to capacity exploitation. Finally, the paper will convey the innovation the industry is engaged in – innovation that moves technology forward globally, and that reaches beyond cable.

INTRODUCTION

Cable operators have seen downstream bandwidth grow at 50% per year (CAGR) for many consecutive years. The trend, often referred to as Nielsen's Law, has held firm for the 20+ years and will be assumed to be as a relevant guideline for assessing the future, along with variants we shall discuss. There are reasonable arguments for long-term limits of media consumption [2,7] that we will consider, although predicting applications has been difficult, and services not yet foreseen may keep the trend alive beyond media consumption.

Cable operators manage this persistent growth under the spectrum constraints of their current legacy service offerings, mostly video, which consume the vast majority of the total available spectrum. Tools for improved bandwidth efficiency are used to balance the growth of legacy services such as HD and VOD as data traffic is increased. Tools and strategies are outlined in [4,10].

Recently, the industry initiated the DOCSIS 3.1 effort, which has an objective to achieve at least 10 Gbps of downstream and 1 Gbps of upstream. This is another major tool for enabling this continued growth, and places cable on par with PON targets, while network migration steps can deliver similar average user capacity.

In this paper, we will take a look at the service growth challenge with an analysis tool concept designed to quantify the problem, introduce and describe in detail the architecture and technology evolutions in play to handle projected requirements, and then revisit our analysis to assess what these can accomplish against this growth.

THE CAPACITY MANAGEMENT TIMELINE

A sample analysis representative of the issue facing MSOs can be charted on a Capacity Management Timeline as shown in Figure 1.

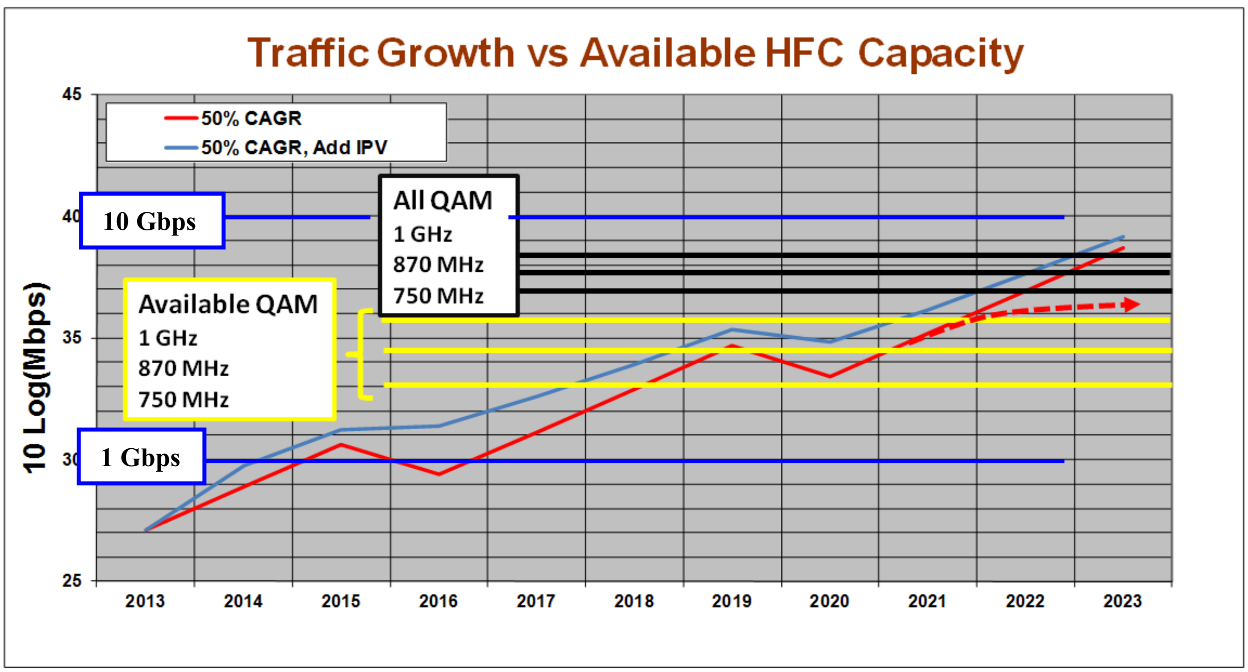

Figure 1 – A Capacity Management Timeline Guides Service and Architecture Evolution

Figure 1 shows various threshold lines drawn that represent the point at which capacity of that particular configuration quantified by the threshold line is exhausted. The purpose of this paper is look at the technology and techniques available that move such thresholds higher to allow more growth, and consider elements that are favorable from a lifespan point of view that effect the trajectories themselves.

So that we can fully appreciate the information in Figure 1 for later use, we will briefly detail the concept of the Capacity Management Timeline. This visual analysis approach allows operators to understand the timing implications of Compound Annual Growth Rate (CAGR) and service evolution. Understanding what it portrays is necessary to make a comparison of the before-and-after of the topics discussed throughout the paper.

The Intersection of Traffic, Services and Architecture

The growth of IP data (DOCSIS) is shown by the red and blue trajectories trending upward with a slope that represents the 50% CAGR. These upward bound trajectories are broken at particular years that represent service group splits (node splits). The blue trajectory has an underlying 50% CAGR, but also includes the introduction of new DOCSIS channels specifically set aside for IP Video (IPV).

Various thresholds are drawn horizontally representing capacity limitations set by the entire forward band using 256-QAM (in black), and the same spectrum examples but offset by channel slots "not available" for growth (yellow, "Available"). In this case, we limited this to 69 slots that were unavailable today. The was based on 60 analog carriers and 9 additional to account for an 85 MHz mid-split for an assumed upstream expansion that takes place over the ten-year window. Legacy digital services obviously coexist, but the idea here was to capture the offset from the all-digital cases ("black") and the power of the analog reclamation step for comparison among capacity management tools.

Of course, any combination of legacy services that add up to 69 channels consumed would yield the same answer. For example, an all-digital downstream broadcasting 200 SD and 100 HD channels would consume about 60 slots. This is just one example – any combination of services can be analyzed and many have been such as in [4,10]. Many specific customer examples have also been analyzed in this fashion, and contrary to what the yellow thresholds might indicate, operators generally do not have *any* room for DOCSIS growth. Actual thresholds are right on top of the current state of DOCSIS consumption. However, this discussion is about new capacity methods more so than bandwidth management [4,10]. We will focus this discussion more on how far "North" we can move capacity thresholds on Figure 1.

Let's take a snapshot of the "today" state from Figure 1 so we can assess the gains we make with the various next generation tools. The aggressive growth of traffic versus time when evaluating against the spectrum constraints looks threatening for HFC's sustainability. The vertical axis is a logarithmic representation to effectively capture compounding growth. Thus, 30 dB represents 1 Gbps and 40 dB represents 10 Gbps. Whenever a trajectory crosses a threshold, that threshold has run out of capacity. For the two cases shown here, the best case scenario with two splits (timed differently than shown for some spectrum cases) manages through 6-8.5 years of IP data growth, without deploying other tools to manage spectrum.

Analog reclamation, Switched Digital Video (SDV), more efficient video encoding, and IP video are all potential tools to help manage the available capacity for growth. The customized use of the Capacity Management Timeline is precisely for this purpose – based on an individual operator's legacy services, growth expectations (data and video), and architecture variables, it is possible to chart out a migration path that allows operators to project their investment needs and timing.

IP Video Transition

There are two growth trajectories on the curve, and these represent a couple of ways to think about and quantify the transition to IP Video. First, note that IP Video will initially be a simulcast, and remain so for many years. Legacy services will co-exist as the video line-up transitions to availability over the IP network. This creates the so-called "simulcast bandwidth bubble," whereby the end state of bandwidth consumption may have an excellent outlook, but the path to getting there is limited by effectively redundant programming.

The two trajectories represent these two views:

1) IP growth at a CAGR of 50% continues to occur, and then on top of that we must add more DOCSIS channels for the IP Video service

2) IP Growth at 50% CAGR has been driven by streaming video services like Netflix for the past several years (not conjecture). 50% CAGR continues because content that use to be elsewhere now joins the IP realm. In this view 50% captures the IP video transition already – this is just the new CAGR growth "engine."

After enough years, as can be seen, the difference becomes very small because the spectrum size needed for IPV is fixed and eventually is overwhelmed by persistent, aggressive CAGR.

The number of IP Video channels required can be determined by analyzing the serving group size, programming line-up, and encoded video bit rates, and understanding the use dynamics of primary screen, secondary screen, and VOD viewing. Also key is a statistical understanding of viewership learned from years of IPTV and SDV deployments. An analysis tool has been developed that does this calculation, and which is publicly available at www.motorola.com/multicast-unicast-calculator/.

A sample case was run with a large SD and HD programming line-up and high penetration of DOCSIS service (70%). The output is shown in Table 1 below. After two splits, about 20 DOCSIS channels are required to meet the IP Video needs (or at one split and 50% penetration early). This is what is added to the 50% CAGR for the blue trajectory, and it was added as 4+8+8 channels over a period of 7 years.

Growth Contraction?

There in one other line of thought regarding the 50% CAGR growth by recognizing that it is being driven by streaming video. This line of thinking is that video quality only increases to a point at which there is no value to improving it [2,7] from a human perception standpoint. It is not completely settled science when that is, but pretty settled that it is finite. The notion that an asymptote exists out in the future associated with video/data consumption (only) is shown by the dashed red line beginning in the year 2021 in Figure 1.

Table 1 – Calculating How Many DOCSIS 3.0 Video Channels

Total DOCSIS Capacity Required to Deliver Managed IP Video Services

Total Bandwidth Requirements - Case 2

Service Group Size HHP:	125	Peak IP Video Devices / Home:	2.5 Active
Max Video Penetration:	70%	At Peak Penetration:	87.5 Homes
Percentage of Video Devices Active at Peak:	70%	At Peak Viewership:	153 Streams

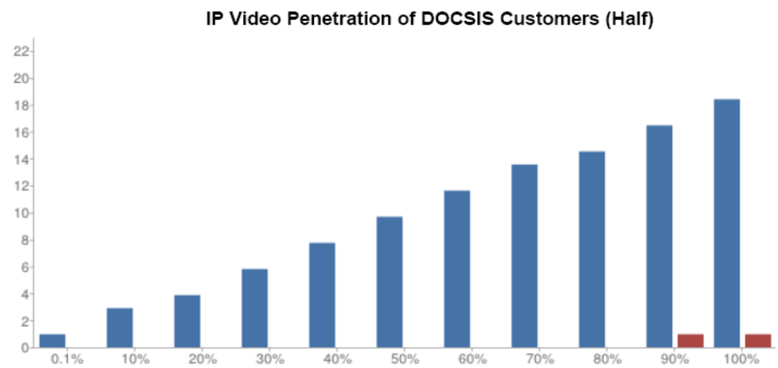

IP Video Penetration of DOCSIS Customers Charts

Total Downstream Channels @ 100% Unicast

Potential Channel Savings with Multicast

There are three principles to this perspective:

1) Assuming media consumption driven bandwidth, we can quantify maximum video quality bit rates that have service value.

2) Recognition that humans have a limited ability to multi-task, in particular with video. While simultaneous secondary screens during a primary viewing may be common, humans have limited ability to focus on multiple things at once with comprehension.

3) Use of IP devices/home and tied to residential demographics which are generally available statistics.

We can also reason that the CAGR engine has been steady for 20 years simply to keep up with increasingly higher levels of human media experiences:

1) Alphanumeric characters
2) Voice
3) Images (pictures)
4) Music
5) Low speed video
6) SD Video
7) HD 1.0

The suggestion here is that perhaps the speeds supportive of the best video quality likely to be practical represent a logical tapering point of CAGR for media consumption as we can fathom it today. There are obviously long-term benefits to HFC networks and migration planning if this does come to pass, as can be concluded by evaluating the implications of the red arrow in Figure 1. We will revisit the implications of this traffic growth philosophy after evaluating our lifespan growth possibilities enabled by new capacity.

Having set the stage for the evaluation of lifespan objectives, let's now look at the component parts designed to paint a prettier picture for that objective, how they do so, and how much they offer.

CAPACITY OPTIMIZATION

Theoretical capacity is a based on two variables – bandwidth (spectrum allocated) and the Signal-to-Noise Ratio (SNR). Shannon Capacity is the well-known limit, and represents the maximum error-free rate that can be achieved in additive white Gaussian noise (AWGN). It is given very simply as

$$C = [B] \text{ Log}_2 [1+\text{SNR (dB)}] \quad (1)$$

This can actually be even further simplified for cable networks, in particular for the downstream, relying on high SNR assumptions. If the SNR is high it can be shown that capacity is essentially directly proportional to bandwidth, B and SNR *expressed in decibels* (dB):

$$C \approx [B] [\text{SNR (dB)}] / 3 \quad (2)$$

This simplification of Shannon Capacity is accurate asymptotically within less than 0.5% with increasing SNR above 15 dB.

Clearly according to (2), more capacity is available with higher SNR, but with logarithmic proportionality. For example, 50% more spectrum yields 50% more capacity, but so does 50% more SNR. However, turning a 30 dB SNR into a 45 dB SNR is a significant network performance leap. Nonetheless, it is certainly the case that more SNR means more capacity, and architectures that create higher SNR – deeper fiber, digital optics, home gateways – create potential capacity.

Shannon capacity is a theoretical concept, and Shannon does not describe either waveform types or codes to use in his famous treatise. For real systems, of course, we deal in signal waveforms and modulation formats to exploit the spectrum. Through this, SNR has two key practical components:

1) Improving the link SNR itself which translates to modulation formats. The link has many contributing noise dependencies – architectural, technology in the optical and RF network, and equipment fidelity and CPE technology itself. The relationship of evolution variables to net SNR impact is a comprehensive accounting of these pieces.

2) Forward Error Correction's (FEC) role in capacity is played out through "SNR" in equations (1) and (2). The best codes enable a given M-QAM format and level of bandwidth efficiency at a lower SNR. Or, alternatively, for a given SNR, the best codes enable the highest order M-QAM formats of the most bandwidth efficiency.

The next section takes a look at the foundational elements of maximizing capacity – optimally exploiting the channel using modern physical layer technology tools.

ADDING TO THE PHYSICAL LAYER TOOLKIT

M-QAM Formats

Today's cable systems implement a maximum M-QAM format of 256-QAM (8 bps/Hz) downstream and 64-QAM (6 bps/Hz) upstream. These represent upgrades in efficiency from prior use of 64-QAM for digital TV downstream and 16-QAM upstream. Through architecture evolutions (deeper fiber) and technology improvements (optical & RF fidelity, DFB return lasers) cable has already gone through at least one major round of improving bandwidth efficiency, and most of it many years ago. Plenty of years have passed since by which a major technology refresh can pay dividends.

Figure 2 shows the current modulation profiles and a couple more that are anticipated as certainties in next generation systems – 1024-QAM and 4096-QAM. "In-between" profiles (512-QAM and 2048-QAM, not shown) are assumed eligible candidates as well. In the figure, all of the M-QAM formats are shown for an *equivalent uncoded BER* of 1e-8. Since they are 6 dB apart for each step up in density, the SNRs are therefore 28 dB, 34 dB, 40 dB, and 46 dB, for 64-QAM, 256-QAM, 1024-QAM, and 4096-QAM, respectively. At the very least, the latter (46 dB) should give pause to the thought of supporting that M-QAM capability over HFC.

Higher order formats can be constructed and, as we shall see, may be worth considering, but are not shown. They do not exist in simulation tools at this point!

A common end-of-line HFC cascade performance requirement for digital channels is a 42 dB SNR with digital channels typically set 6 dB below analog channels. Given that 256-QAM requires 34 dB (1e-8) without coding, and up to 4 dB less than this by DOCSIS specification with a J.83 Annex B error mitigation subsystem included, it is apparent why today's networks are very successful with 256-QAM. In fact, some are likely able to support 1024-QAM robustly using similar "J.83B" tools. Some lab evaluations have indicated this is likely to be the case [9].

However, even just considering HFC SNRs, the 1e-8 SNRs required of 2048-QAM (43 dB) and 4096-QAM (46 dB) clearly indicate extra "help" is necessary to achieve these with robustness. It can come in the form of FEC, architecture modifications, technology improvements, or all of the above, as long as we can find the dBs necessary to close the link.

A View from the Field

Figure 3 shows some extremely valuable pioneering work done by a major North American MSO – a first of its kind that indicates with a large statistical sample what Cable Modems are telling us their channel SNR looks like [15]. Other MSOs are now gathering such statistics as well to help the industry engage in proper technology choices based on real data.

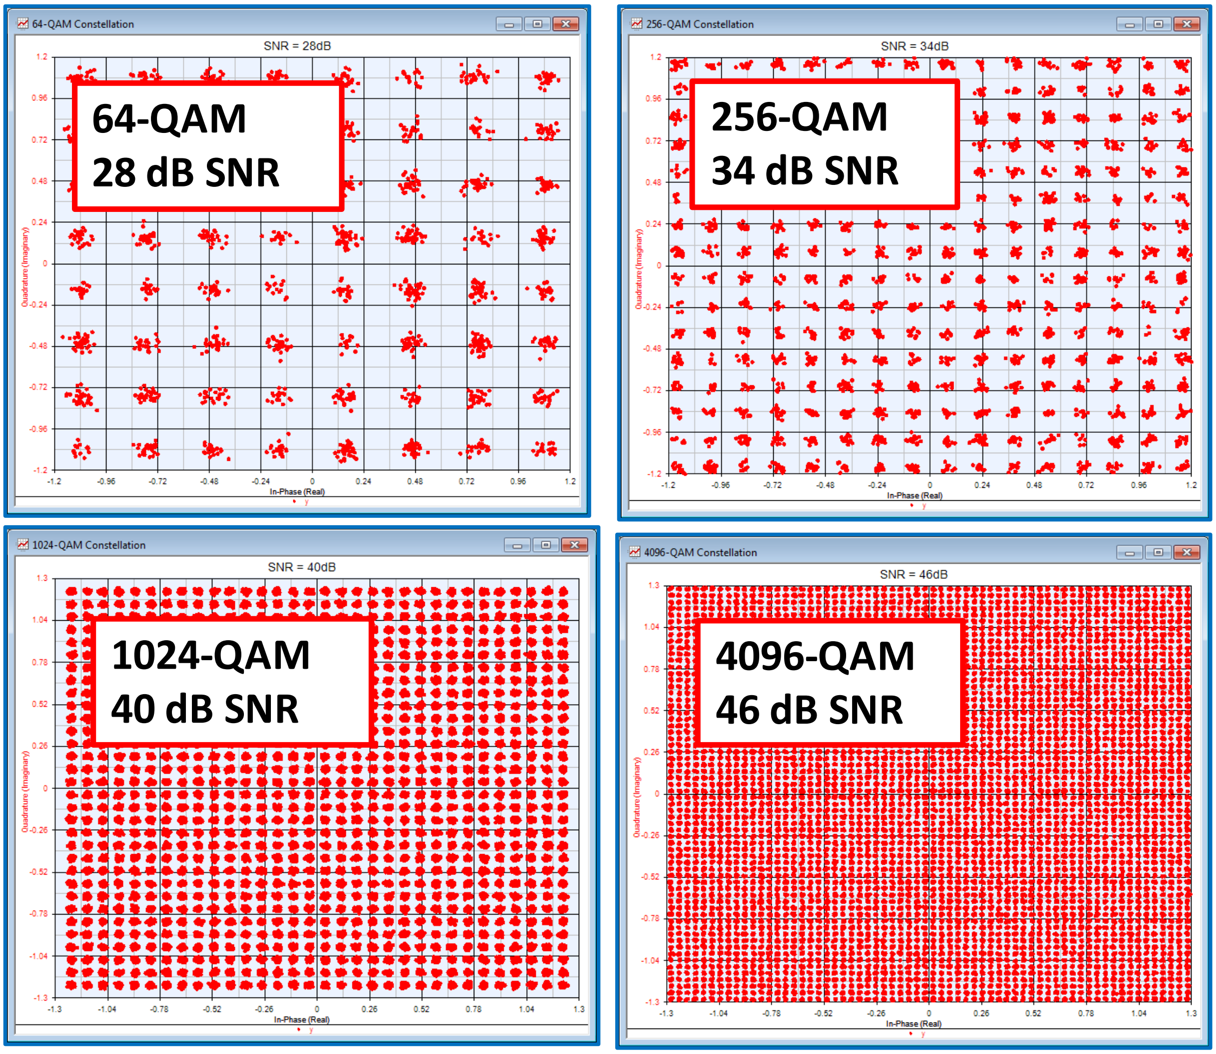

Figure 2 – Increasingly Spectrally Efficient M-QAM Formats

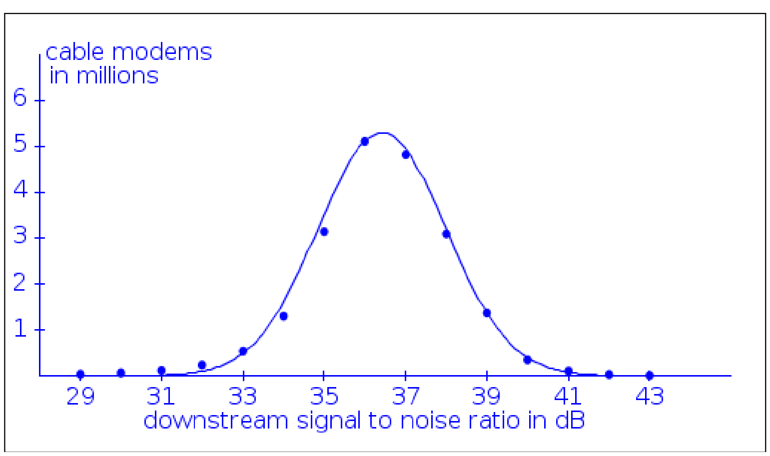

Figure 3 – Major MSO Cable Modem SNR Distribution [15]

There are important differences between CM reported SNR and HFC delivered SNR, as we can easily determine by the delta between the HFC delivered 42 dB number (or better) and the SNR scale in Figure 3. The most important ones are:

1) The CM actually measures and reports MER, which includes all impairments on the channel, all the way to the CM demodulator. Thus, it includes the CM contribution itself.
2) The CM's contribution is strongly dependent on the location of the CPE in the home. It is a dominant noise contributor at low CM input levels.
3) The CM was implemented for high performance of 256-QAM, which is 12 dB less sensitive than 4096-QAM.
4) The maximum measurement fidelity itself of MER is likely in the low-to-mid-40's.

Figure 3 will prove valuable in defining QAM formats and techniques to optimize their use. While the absolute SNR numbers may be biased towards lower values relative to a new generation of technology and architecture evolution, the spread of the distribution is illustrative of the variation across the network that can be better exploited for capacity management.

<u>The Magic of FEC</u>

Advances in FEC have straightforward PHY design effects – better FEC reduces the SNR required to achieve a particular QAM format, increasing bandwidth efficiency and throughput for a given link performance. Today's go-to code family is Low Density Parity Check Codes (LDPC). LDPC codes have been mathematically around for many years. However, as has been the case with other codes (e.g. Reed-Solomon), they have came into vogue as the speed of computation has become sufficient to enable real-time operation of these extremely resource-intensive large block size codes. The first standard to define an LDPC code was DVB-S2 in the early 2000's, but since that time codes from the LDPC family have become part of G.hn, MoCA™, WiMax, Wi-Fi, and DVB-C2, among others. The reason is simple – they get closest to the Shannon bound, maximizing capacity, and efficient ways to implement them cost effectively are now available.

In Figure 4, we show the DVB-C2 family of LDPC codes [14] and the M-QAM potential available, including 64-QAM through 4096-QAM. Observe the SNR requirements enabled by LDPC under the "Highest Code Rates" label in Figure 4 (90%). These are the nearest apples-to-apples comparisons to the error correction scheme used by J.83B downstream today.

The true power of LDPC can be seen in the SNRs required to deliver vanishingly low error rates in Figure 4 and Table 2. Table 2 summarizes the SNR gains available for the QAM profiles compared to the uncoded case [6]. The FEC, of course, comes with a 10% efficiency penalty (for the 90% code rate). However, 10% efficiency hit for *9-11 dB* of SNR gain is a powerful trade-off – *one-tenth* the SNR tolerated for this small loss of efficiency. The 46 dB of uncoded 4096-QAM SNR previously mentioned, for example, reduces to 35 dB as shown in Figure 4 – pretty impressive! The 9-11 dB range of SNR advantage in Table 2 is a testament to the power of LDPC codes. We will compare this to today's downstream FEC in the next section.

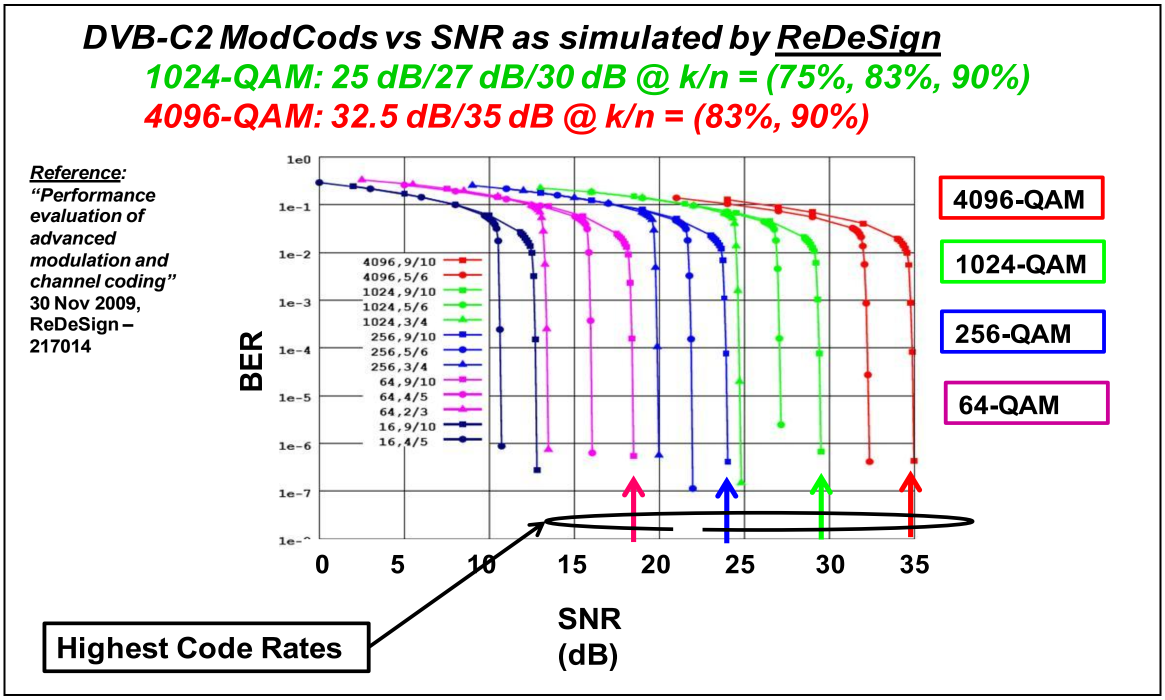

Figure 4 – Bandwidth Efficient M-QAM Enabled by LDPC FEC

As impressive as Table 2 may look, M-QAM constellation pictures truly put the role of FEC into perspective. To emphasize this "magic," we show the constellations of 1024-QAM and 4096-QAM in Figure 5. The SNRs shown are 3-4 dB *higher* than the SNR threshold for low error rate (error free) performance in Figure 4. Figure 5 is the "picture is worth a thousand words" version of Table 2, illustrating the power of FEC to clean up what is quite an incoming mess.

Table 2– Coding Gain of LDPC FEC

	Uncoded ~1e-8	LDPC DVB-C2 @90%	SNR "Gain" (dB)
64-QAM	28	19	9
256-QAM	34	24	10
1024-QAM	40	30	10
4096-QAM	46	35	11

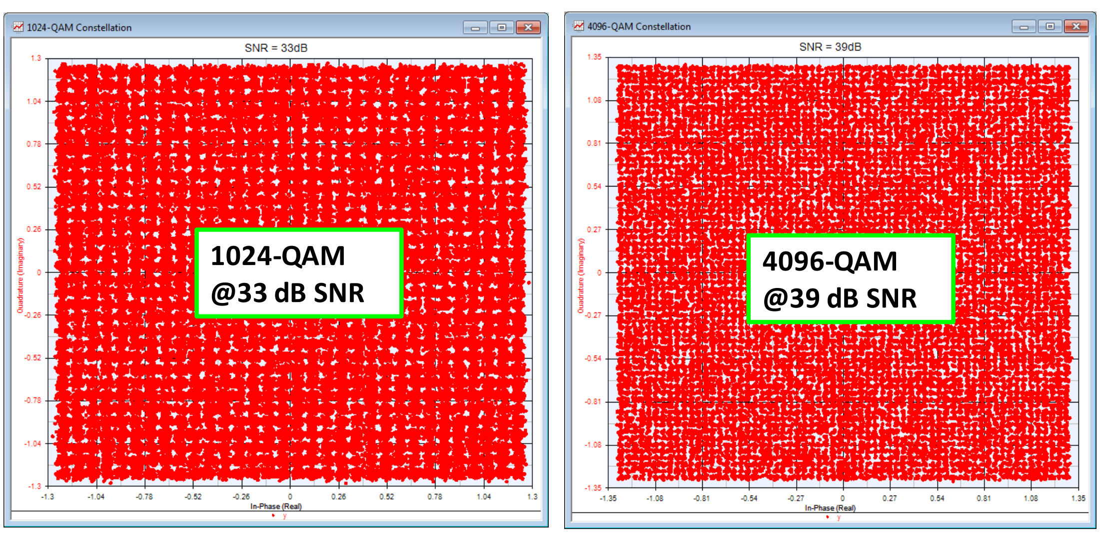

Figure 5 – Amazingly Error Free: The Power of LDPC Forward Error Correction

FEC II – How Does it Do That?

We can precisely identify the dB of FEC advantage of LDPC versus today's ITU J.83 Annex B downstream (as well as for the upstream). Refer to Figure 6 [1].

In Figure 6 (simulations by Intel), we can compare SNR vs. Code Rate for the old and new FEC choices. For the downstream, J.83B (orange) can be compared against the DVB-C2 short (red) and long (blue) codeword. The plot is based on 256-QAM, with the expectation that similar relationships will hold for other M-QAM formats for a well-designed code. Note that J.83 Annex B does not actually have variable code rate, but varying the Reed-Solomon code rate enables a relevant and straightforward simulation while allowing apples-to-apples code rate comparisons.

Figure 6 identifies how, with LDPC alone, we could actually manage a two-order increase in modulation profile – a 6 dB theoretical SNR gap – using a combination of the code family and code rate, if this were desirable, as follows:

- Labeled by the orange crosshair and bracket, LDPC at the same code rate provides about 3.2 dB of SNR gain (red bracket) compared to J.83B. A 3 dB change is roughly the equivalent of one half-step modulation order, such as 256-QAM to 512-QAM.

- At the cost of efficiency, by reducing the code rate by about 10% (to 80%), another 2.7 dB can be gained, for a total of 5.9 dB, or nearly 6 dB (green bracket and horizontal arrow).

Thus, a little more than 3 dB comes from the change in code family, and the rest comes from a 10% drop in the code rate. Since the code rate is an efficiency reduction, some or the rest of the difference to get to a 6 dB difference, such as 256-QAM to 1024-QAM, might instead be made up, for example, with architecture or technology evolution in the HFC network.

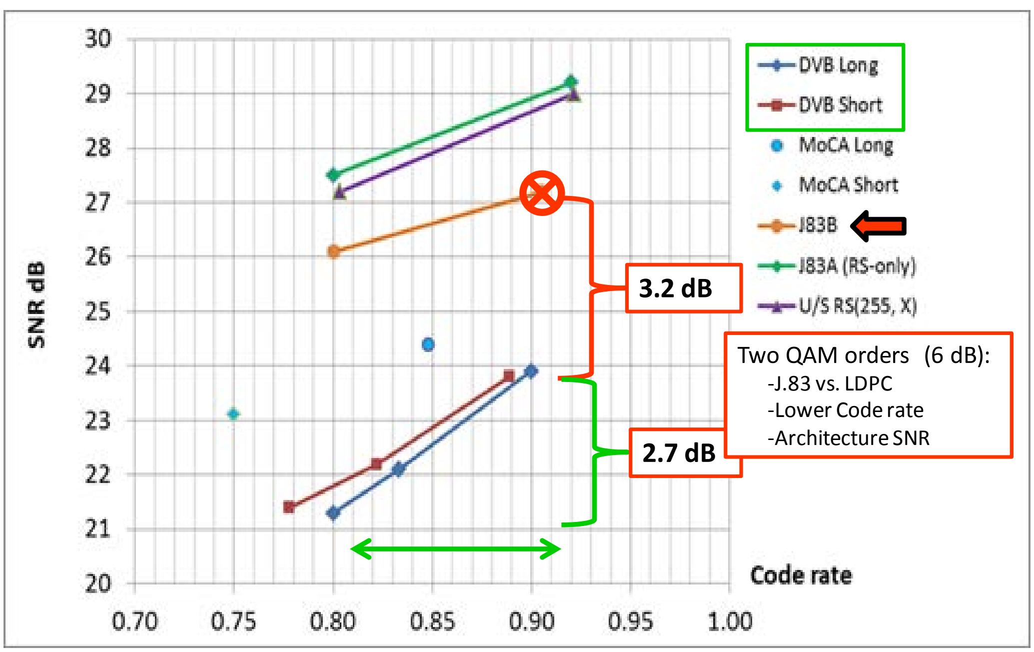

Reference – Mission is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS, 2012 Cable Show Spring Technical Forum

Figure 6 – LDPC vs. J.83 Annex B Comparison (Downstream) [1]

We can perform the same analysis for the upstream, as shown in Figure 7 [1]. Today's upstream *does* have a selectable code rate. The cases for t=10 and t=16 symbol-correcting are shown in the simulation (courtesy of Intel). We show to two MoCA™ codes and compare to the MoCA™ short code. The availability of shorter codeword sizes is essential to match the upstream packet size distribution.

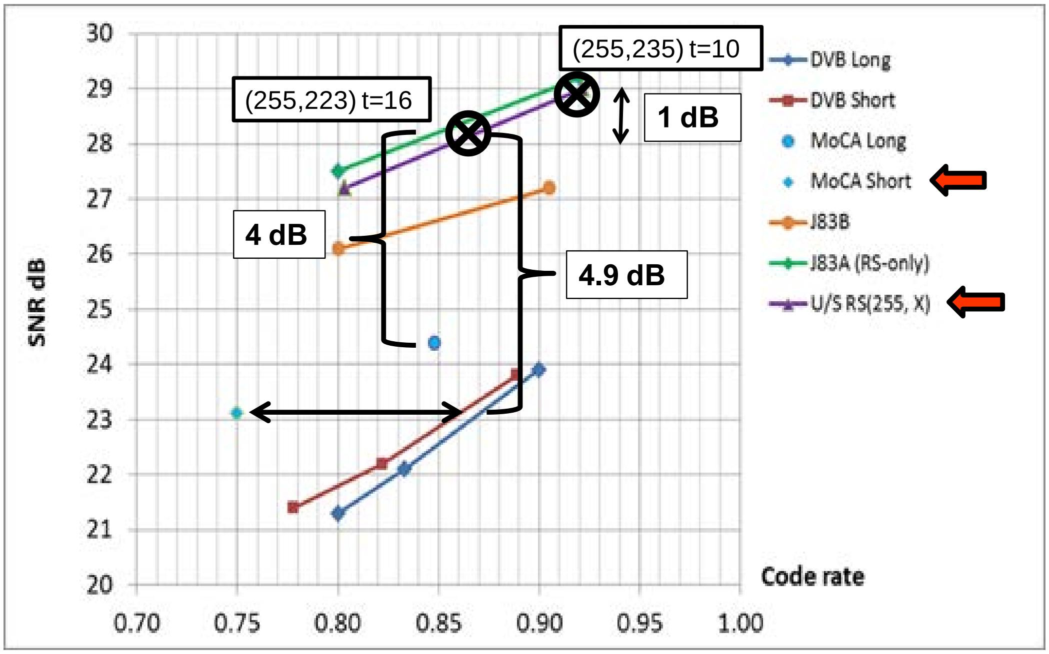

Figure 7 – LDPC vs. Reed-Solomon Upstream Comparison (Upstream) [1]

As Figure 7 shows, we can again work out the potential for a two modulation order improvement. Using the MoCA™ short code (blue diamond), we note that the SNR requirement is (4.9 + 1) = 5.9 dB lower than the t=10 error correcting. This comes at the cost of lower code rate (by 17% - significant) and thus lost efficiency. The efficiency loss when comparing the MoCA long code to the t=16 case is much less (2%), but we do not achieve 6 dB, only 4 dB. However, we might consider upstream technology or architecture improvement that offer 2 dB of additional SNR link budget to close the gap.

Since the upstream optical technology tends to be the dominant factor in the upstream SNR, the ability to directly affect the upstream bandwidth efficiency is more straightforward than the downstream. Headend de-combining is another area where instantly accessible dB can affect the upstream bandwidth efficiency potential.

M-QAM, FEC, SNR: Connecting the Dots

With knowledge of both lower M-QAM thresholds enabled by LDPC FEC, and a well-quantified awareness of the SNR on the receiving end by fielded cable modems, we can connect the dots between the two to examine the potential for new downstream capacity. Figure 8 shows the two together to begin this comparison [16].

The Figure 3 distribution on the lower right – a classic Gaussian bell curve – shows an average of about 36.5 dB and a 2σ variation of about 3 dB. This puts over 95% of the measured modems from this large sample between 33.5 dB and 39.5 dB (±2σ).

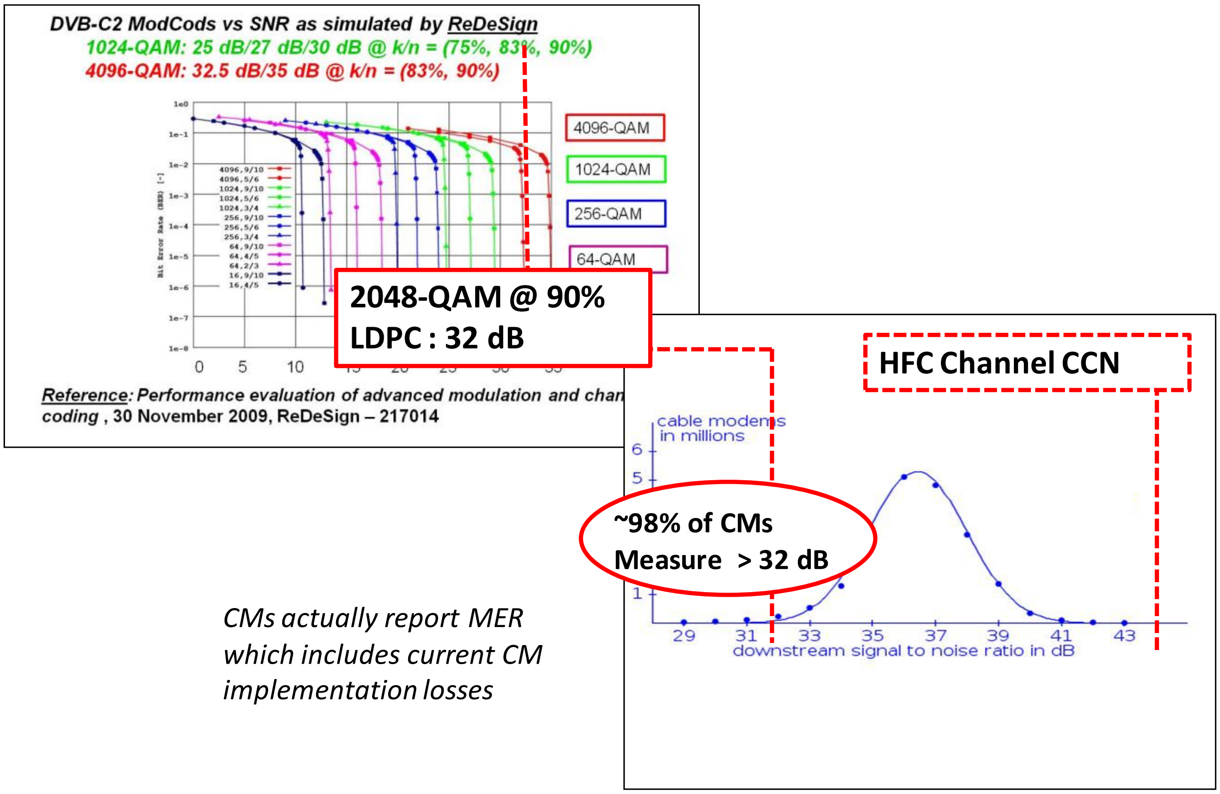

Reference: http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf

Figure 8 – M-QAM Potential Based on Today's Measured MER Characteristics [16]

The Figure 4 QAM-FEC simulations repeated in Figure 8 do not include the mid-step constellations. However, they are easily estimated, and in this case the estimate for 2048-QAM for the 90% code rate would be that it is 3 dB lower than the 4096-QAM SNR requirement of 35 dB, or 32 dB. On the CM distribution curve, this represents a performance achieved by about *98%* of the modems. This shows, not surprisingly, that using only 256-QAM leaves potential capacity on the table. Note that 256-QAM @ 90% DVB-C2 LDPC requires a 24 dB SNR, which only re-emphasizes the point.

Of course, this does not account for added operator margin required for robustness. A substantial margin is used by field technicians to guarantee a robust 256-QAM downstream today. Figure 6 shows the 27.5 dB of SNR required for 256-QAM in the J.83B downstream. Typically, operators will look to obtain about 35 dB (operator dependent) to "certify" an install at a customer's home [18]. We will address the margin topic more specifically in a subsequent discussion about downstream optimization, as we anticipate that this paradigm will change. However, for now, we can recognize that using a 2σ spread's lower SNR edge of about 33.5 dB in Figure 8, and subtracting the equivalent 7.5 dB margin we are left with 26 dB as an SNR. Based on Figure 4 would support 1024-QAM with a code rate close to 80%.

Lastly, note the "HFC Channel CCN" label and red line on the lower right of Figure 8. CCN stands for Composite Carrier to Noise, accounting for both AWGN and digital distortion build up which looks like AWGN from a noise floor perspective. It is the HFC plant equivalent of SNR. This line describes what the plant can deliver at end of line (EOL). Minimum performance of 42 dB has previously been mentioned, while typical performance is higher such as that shown here. The point here is that the HFC channel, if properly implemented, is *not* limiting capacity from an SNR (CCN) perspective.

In summary, it should be obvious that 256-QAM is not the best case bandwidth efficiency possible in the downstream. More bps/Hz are available if we desire to chase after them. Moreover, some of the most important capability to obtain these bits is already in place, in particular around the HFC channel quality as is understood in terms of minimum EOL today, and even as reported by the CM SNR data in Figure 3 which accounts for a broader set of variables which will only improve with architecture and technology evolution. Therefore, if we need more bits, they are not far from reach. And, as Figure 1 implies, the question "do we need them" has already been answered.

The Role of OFDM

An element hidden by the capacity equations in (1) and (2) is the accuracy of a constant, static, and spectrally flat assumption of SNR. In many systems today – particularly wireless – the SNR can be quite dynamic when moving throughout a cell, for example For other channels, such as cable, it is not particularly dynamic, but does vary across the area it serves both geographically and with respect to frequency of operation.

Also, the frequency response of the channel has large implications on the receiver design and its ability to perform close to the spectral efficiency that the channel SNR suggests it should achieve. For wireless, moving across a cell in a metro area creates a difficult multi-path environment. In cable channels, a wide range of ripple and slope may exist due to static channel multi-path (micro-reflections in cable-speak) conditions as well as due to the nature of having a multi-octave RF distribution network and serving uncontrolled home coaxial architectures.

Variable and unpredictable channel conditions are specifically where multi-carrier systems (e.g. OFDM) come into play. The fundamental OFDM concept is shown in Figures 9-11.

The fundamentally different characteristic of OFDM is replacing classic single-carrier QAM, such as the 6 MHz and 8 MHz channels used today for nearly all QAM signals on the cable, with many narrower, subcarriers, and sending these subcarriers in parallel. This is depicted in Figure 9. Narrow means kilohertz-type of narrow. As a practical example, 10 kHz subcarriers would mean there are 600 of them inside a 6 MHz "normal" channel slot in North America. As in single carrier technology, the subcarriers themselves carry QAM, which is why we study QAM modulation formats in detail regardless of RF waveform type. In the ideal AWGN environment, the two techniques perform equivalently.

The other uniquely interesting OFDM characteristic is that the narrow subcarriers overlap by design, as shown in Figure 9. They get away with this (clearly classic frequency division multiplexing, or FDM, could not) by maintaining a relationship among subcarriers that connects their spacing to the symbol rate so that they remain orthogonal. Ideally, orthogonality ensures that, by the nature of the waveform integration during demodulation, subcarriers do not interfere with one another.

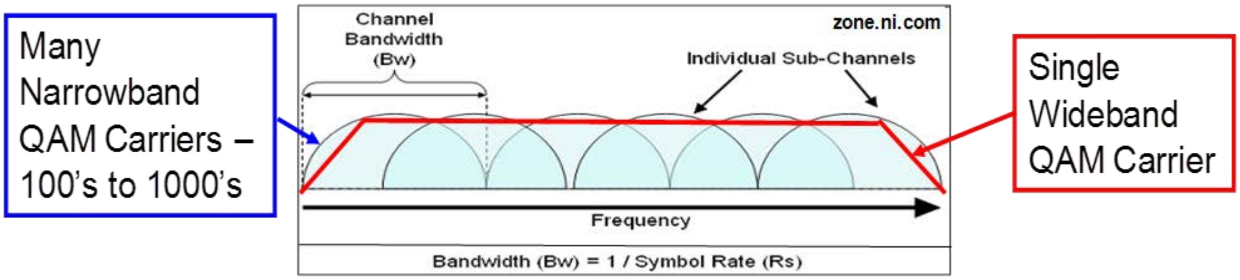

Figure 9 – The Multicarrier (OFDM) Concept: Frequency Domain [20]

In the time domain, this "zero interference" quality is shown in Figure 10 whereby integrating (detection) over the period shown for one of the subcarriers has the others summing to zero. Figure 11 shows the frequency and time aspects together. All subcarriers are sent in parallel during a symbol transmission, and the process is repeated at the next OFDM "symbol" transmission.

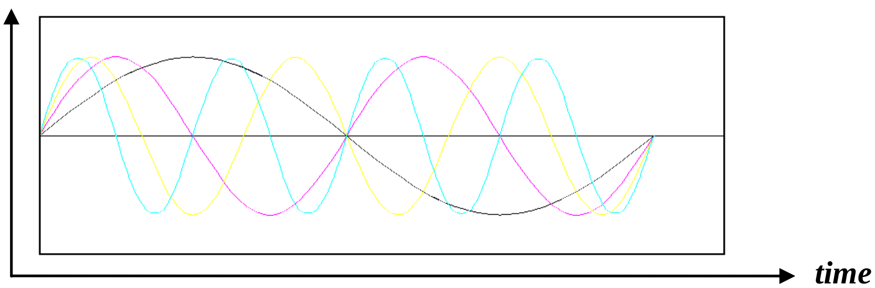

Figure 10 – The Multicarrier (OFDM) Concept: Orthogonality in the Time Domain [13]

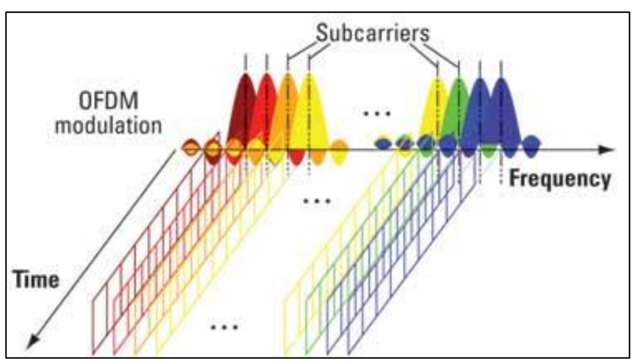

Figure 11 – OFDM: Frequency and Time Domain [19]

The next transmission does not immediately follow the first (at least in terms of payload transmission) – this is one of the fine details of OFDM system design we will not get into here, but which deals with how OFDM effectively performs the function of "equalization." OFDM uses what is called a "cyclic prefix" (CP) to delay a new data transmission beyond the multipath window.

The whole OFDM idea sounds unnecessarily complex, and indeed this was once the case. Like FEC, the multi-carrier concept was invented by brilliant engineers who noted many of the potential benefits reaped from this approach to accessing a channel many years before the implementation became practical. We will not get into implementation details, but OFDM was largely made practical, and actually quite simple, with advances in real-time computing power than enabled wideband, high-speed, high resolution FFTs that could be processed in real-time.

Shannonizing with OFDM

A good way to interpret the OFDM approach in terms of its capacity-maximizing effect is to write the expression for capacity in (2) in "long" form:

$$C \approx (1/3)\sum_{\Delta f}[\Delta f]\ [P(\Delta f)\ H(\Delta f)\ /\ N(\Delta f)]_{dB} \quad (3)$$

Here, instead of bandwidth, we have used a summation of spectrum chunks using a set of small frequency increments, Δf. The sum of all Δf increments is the bandwidth available, B. Instead of SNR, we have broken it down into its components: signal power (P), noise power (N), and channel response (H) – each also over small Δf increments. In practice each Δf represents the width of one OFDM subcarrier.

The total capacity above is then simply the summation of the individual capacities of chunks of spectrum. The purpose of the form used in (3) is to recognize that channels may not have a fixed SNR characteristic, such as due to expected non-flat frequency response variations and uncharacterized spectrum above today's 1 GHz forward band. In this case, the capacity of a not-flat SNR region can be calculated by looking at it in small chunks that, because of their narrow width, themselves approximate flat channels. A similar argument applies when there is, for example, interference. The effected OFDM subchannels will have a lower SNR (in this case S/(N+I). This flexibility is a key advantage of multi-carrier modulation such as OFDM – very narrow channels, each of which can be individually optimized.

For a single carrier transmission, it becomes increasingly difficult for wider and wider channels to achieve the same effect without complex, and sometimes impractical equalization techniques and interference mitigation mechanisms. Or, in the case of DOCSIS, it becomes impractical to channel-bond more and more single-carrier channels without incurring excessive complexity and inefficiency.

The long-form capacity equation above demonstrates why OFDM is often better suited to achieving the best throughput possible than single-carrier techniques in channels with poor or unknown frequency response, and in particular when that response is time-varying.

The HFC downstream is typically very high SNR and generally well-behaved. However, it can be subject to large broadband frequency response variations when signal reflections are high. The downstream is also increasingly susceptible to 4G interference as these deployments increase, as well as interference sources that have existed for years. Outside the current downstream – above 1 GHz – plants are likely to vary widely as there are no requirements to be met or equipment specifications that can be used to help define the spectrum, though the coaxial cable medium clearly can be exploited beyond 1 GHz.

In the upstream, the channel is much less predictable than the downstream, particularly at the low end of the band, and burst noise events are more prevalent than in the downstream. Furthermore, the upstream is as likely if not more so than the downstream to see a bandwidth extension into new territory, such as 85 MHz and even to 200 MHz. However, because of its "funneling" architecture, interference that may be localized and insignificant in the downstream today may impact the channel for all in the upstream when the diplex is adjusted for more upstream spectrum. The FM radio band is the most obvious candidate should the upstream extend beyond 85 MHz. The interference-protection properties of OFDM will be valuable in this case, as it is in the troubled part of the return band today.

Note that in the upstream, the likely multi-carrier candidate is actually OFDMA, or Orthogonal Frequency Division Multiple Access. The principles of the signal waveform are the same, but in the case of OFDMA, different sub-channels can be allocated to different users simultaneously, an attribute important to efficient use of the upstream. The difference between OFDM and OFDMA is shown in Figure 12. We will generally use just "OFDM" to refer to the technology in both upstream and downstream.

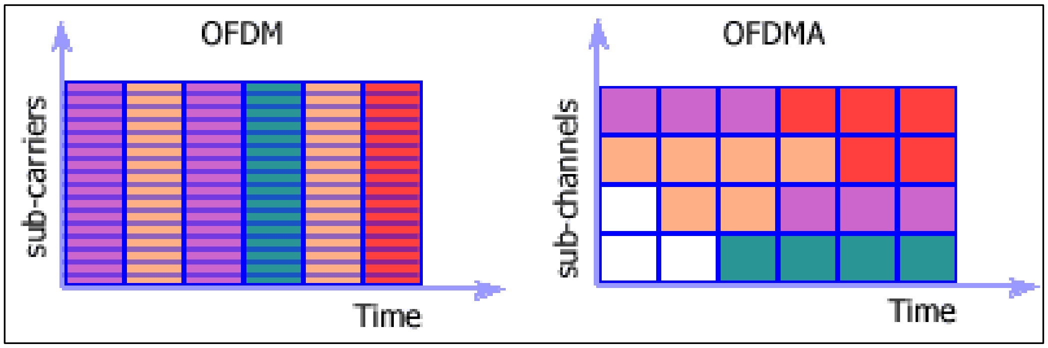

Figure 12 – OFDM vs. OFDMA [12]

As discussed previously, supporting more bandwidth efficient M-QAM profiles over HFC has little to do with whether we are discussing single carrier QAM or OFDM-QAM. When it comes to SNR (AWGN), system performance is identical. OFDM's most valuable HFC role is to overcome frequency response characteristics and unknown channel quality and manage interference conditions to yield the best probability of maximum SNR exploitation for capacity. Very wideband (high-speed) operation is also a major plus. Historically, OFDM applications have been linked by this common thread – unknown or poor RF channels – and the benefits it provides in those cases are being brought to the cable environment. In the downstream, the most questionable spectrum would be the band above 1 GHz, and in the upstream the entire channel is more suspect, but especially so 5-20 MHz.

Relative to bandwidth above 1 GHz, Figure 13 [1], shows the range of insertion loss characteristics of various models of a *single tap* type above 1 GHz for "1 GHz" specified taps. It is clear that any given tap, much less a cascade of taps, will be highly unpredictable from system to system, and even from RF leg to RF leg in the same system.

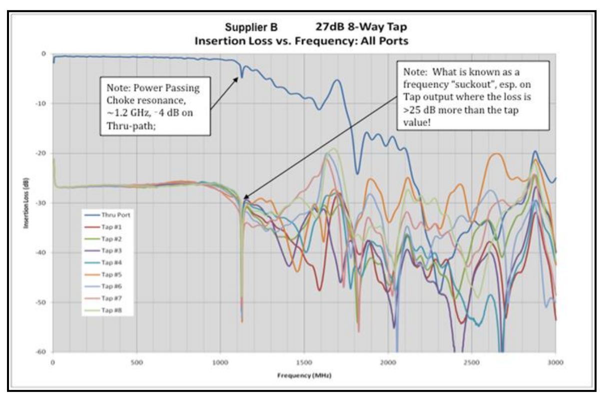

Figure 13 – Unpredictable Frequency Above 1 GHz [1]

There are other important OFDM benefits not associated with system performance. Some of these are listed in Table 3.

The second point in Table 3 is perhaps the next strongest argument for OFDM for

HFC, albeit it a more practical "operations" one. With so much spectrum and service evolution anticipated over the next decade, the granular spectrum management enabled by OFDM through flexible subcarrier allocation (using some but not all subcarriers) is a valuable tool when working around a full band of legacy spectrum.

Table 3 – Why Cable OFDM?

Why OFDM?
Optimizes Channel Capacity, in particular for unknown, uncharacterized, and hostile interference channels
Granular spectrum allocation beneficial during band plan and service transitions
Multiple sources of supply and likely cable investment
Consistency with **other standards** and cable network extensions (wireless, EPoC)
OFDM + LDPC to Layer 1 as IP is to Layer 3 – likely final RF step (little more capacity worth exploiting)
Implementation complexity favors OFDM over TDM afor wideband channels with linearity distortions
More **Spectrally Efficient** Wideband Channel than NxFDM, 2-D Multiple Access (OFDMA)

Other points in Table 3 worth mentioning include the increasing ability to do computationally complex operations in real time. OFDM implementation – once the major obstacle – has become a strength through IFFT/FFT functionality that forms the core of transmits and receive operations.

This implementation advantage leads to one of the final strong, business-oriented arguments for OFDM. As an ecosystem, the number of suppliers of OFDM technology and the range of industries engaged in it enlarges the pool of technology resources and leverage tremendous economies of scale. The wireless industry and Home LAN products in particular both represent very high volume applications.

Impairments: Single Carrier and OFDM

OFDM puts a different signal type on the wire, and because of that it responds differently to some of the common impairments of cable – unique (CTB/CSO) or otherwise (additive interference, phase noise). We mention these two important ones here, but for a fuller treatment refer to [6]. An understanding of the differences will be critical to properly specifying and operating OFDM on the cable channel, and analysis of these effects in ongoing.

CW Interference

Single carrier techniques combat narrowband interference through adaptive filtering and equalization mechanisms. OFDM, on the other hand, deals with narrowband interference by avoidance. Also, what may be "narrow" for a single carrier QAM signal may not be narrow relative to an OFDM subcarrier. Figure 14 shows OFDM impinged upon by two interferer types – a CW carrier and a modulated waveform of some unspecified bandwidth but that is similar to OFDM subcarrier spacing.

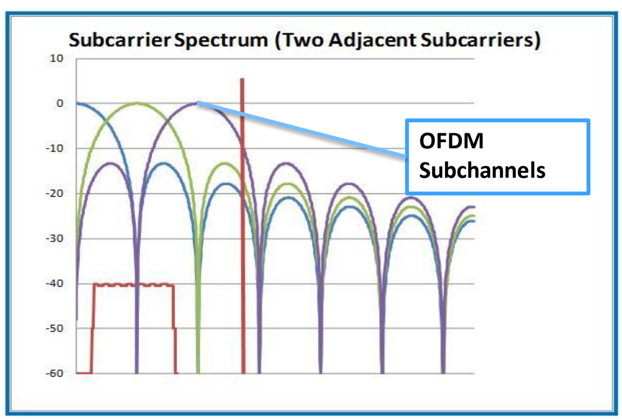

Figure 14 – Interference as Seen by OFDM

Subcarriers imposed upon by an interferer can be nulled or modulated with a more robust modulation profile. The effect is a capacity loss, but generally a modest one because only a limited number are affected. Compared to SC-QAM, OFDM offers graceful degradation via lost capacity, as opposed to a thresholding effect at some intolerable level of interference. This could be viewed as both pro and a con. SC-QAM, for example, may find low levels of interference essentially invisible from a detection perspective, a scenario well represented by analog CSO/CTB distortion beats in the forward path.

CTB and CSO when analog video is present also have more of a deterministic quality – always preferred – in location, level, and whether they will even be present or not. Figure 15 compares 6 MHz SC-QAM and OFDM-QAM with respect to CTB/CSO interferers.

Two key characteristics stand out:

1) Distortion beats are no longer necessarily narrow relative to the subcarrier bandwidth, on average. The distortion bandwidth and amplitude vary slowly, however, and these peaking effects can have well-documented implications for QAM performance and interleaver depth.

2) Beat amplitude is much higher relative to SC-QAM since each subchannel is a small fraction of the total signal power in, for example, 6 MHz. For the 600 subcarriers per 6 MHz example, this is 27 dB. So, CTB/CSO of 53 dBc is now 25 dBc! And, that is just the average, not including its amplitude modulation characteristics. Clearly, for OFDM, the FEC will be required to deliver error-free bandwidth efficiency.

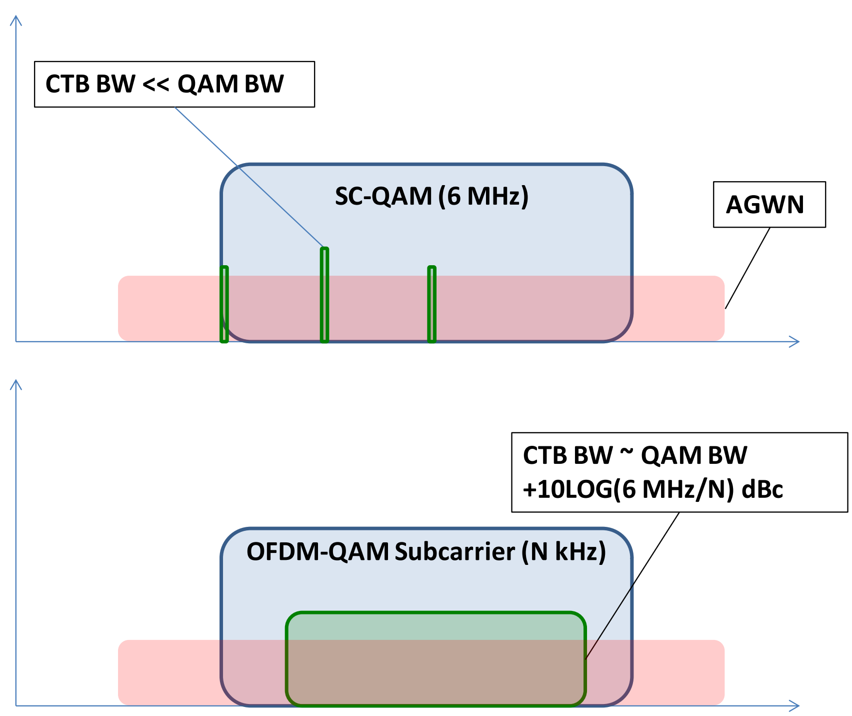

Figure 15 – CTB/CSO Interference – SC-QAM vs. OFDM-QAM

OFDM system design and choice of parameters for the error mitigation subsystem are used to overcome interference in the channel whether the mechanism is distortion beats or additive interference. The latter is being observed in some cable systems in LTE bands.

Phase Noise

OFDM creates an interesting scenario with respect to phase noise degradation. A typical assumption for SC-QAM is "slow" phase noise. The exact spectral mask is less important – only the untracked rms phase noise matters. For OFDM, with many narrow sub-carriers, the phase noise mask will typically extend *beyond* the subchannel bandwidth. Figure 16 shows a characteristic lowpass shape of untracked phase noise (two cases of different "bandwidth") against an OFDM sub-channel spectrum.

Phase noise thus includes two degradation mechanisms for OFDM. There is an error common to all subcarriers related to the "in-band" effects and known as "common phase error" or (unfortunately) CPE. This component is often largely tracked out and therefore of little consequence, and has the classic "rotation" effect on each subcarrier (thus "common"). The other typically more impactful component is that associated with Interchannel Interference (ICI) as the mask cross into other sub-channel bands and these effects are summed. This phase noise effect is additive, uncorrelated noise, which is better than rotational CPE, but the ICI effect from adjacent subchannels has the potential to be quite high.

Both CPE and ICI effects must be accounted for in system design, and the techniques for doing so are well understood. There are just significant differences in how to approach the solution compared to the traditional single carrier design, and we will be attempting to do so with much more sensitive, higher bandwidth efficiency, M-QAM formats.

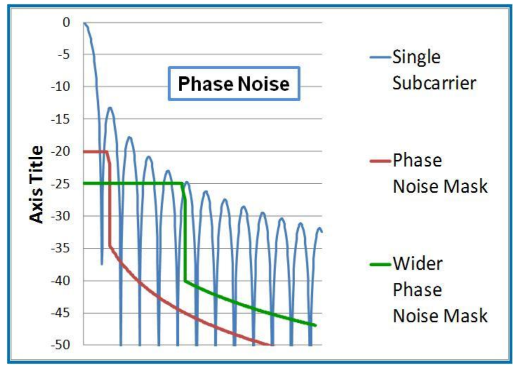

Figure 16 – The Shape of the Phase Noise Mask is Critical for OFDM

Towards a Layer 1 Standard

OFDM offers a robust way to exploit spectrum above 1 GHz, which will be necessary to achieve the objective of 10 Gbps for DOCSIS 3.1. It also provides advantages in the downstream and upstream as interference sources arise going forward, and provides robustness at the low end of the return that can only be managed with S-CDMA today. In addition to its capacity optimizing capability, because of its granularity of spectrum allocation, OFDM provides bandwidth efficient flexibility for systems undergoing service and spectrum evolution which could prove very valuable. With the transition to IP, continual enhancements of HD and on-demand services, managing a full downstream and nearly full upstream, and the expectation of a new diplex crossover sometime in the future, this is a valuable benefit.

In the long-term, OFDM-QAM plus LDPC FEC, because of its capacity optimizing capability in any channel and implementation simplicity, can be viewed for OSI Layer 1 (PHY) what IP has become for OSI Layer 3 – a de-facto go-to standard. This enables potentially significant simplification of network evolution over time. Virtually all modern RF systems across multiple industry segments implement some form of OFDM – 4G Wireless, Wi-Fi, MoCA™, G.hn, HomePlug AV, 802.11n, and VDSL. This end-state scenario would be very similar to what is currently happening in the all-IP transition, where in that case we are simplifying at the network layer. The standardization would extend to the lower layers of the stack and include some components of a software-defined architecture.

♪ ARE YOU READY FOR SOME 4K? ♫

The performance of HFC networks in the downstream is very well understood from decades of achieving fidelity acceptable for analog video. Some typical performance numbers are shown in Table 4 for the case of 60 analog carriers on a 750 MHz system over a range of cascade depths for two different return spectrum scenarios.

These first four columns are referenced to analog levels, so for digital they must be lowered 6 dB. This is listed in the far right column as "QAM CCN". Again, CCN captures all noise floor components – AWGN and digital distortions – and is for all intents and purposes is HFC's SNR. Digital distortion contributors are many and largely independent, so a Gaussian assumption is reasonable.

An important and expected result from the table is the improvement in the CNR and QAM CCN as the cascade shortens and service group size gets smaller. As fiber penetrates deeper, average bandwidth per subscriber is increased, but also the channel quality improves. An RF cascade has the effect of cascading degradation at every amplification stage in the downstream – both noise and distortion. In the upstream this is also the case, but to a lesser degree of importance, while the shrinking service group size has more significant benefits to channel quality, associated primarily with interference funneling.

Table 4 – Downstream Performance vs. Cascade

		750 MHz				
		60 Analog				QAM
		CNR	CSO	CTB	CCN	CCN
Return	N+0	54	64	67	53	47
5-42 MHz	N+3	51	60	63	50	44
	N+6	50	58	60	48	42
	N+0	54	64	67	53	47
5-85 MHz	N+3	51	60	63	50	44
	N+6	50	58	60	49	43

Figure 17 illustrates the fiber deep concept from a cable serving area footprint perspective. Note that service group splitting may also be achieved simply though a segmentable node, with no effect on the cascade depth. In this case, it is the upstream channel primarily that benefits.

Network Nirvana

A somewhat natural architectural end-state vision for HFC is business-as-usual node splitting culminating ultimately in an N+0 system – a passive coaxial last mile with no RF actives after the node. The benefit in terms of channel performance can be observed from Table 4, where now everyone is the “N+0” column.

Besides the channel quality improvement afforded by N+0, a very important advantage of the “fiber deepest” architecture is the ease with which new capacity can be exploited without the existence of actives in the path. Actives involve diplexers, which add obstacles to adjusting spectrum allocations, and their ability to supply quality “excess” bandwidth to 1.2 GHz the way taps may is probably more questionable being active circuits.

It is assumed that HFC link performance can be maintained as the spectrum shifts to higher spectrum in the case of 1.2 GHz. The loading effect of *increased* total spectrum, such as 108 MHz-1200 MHz, can also be calculated.

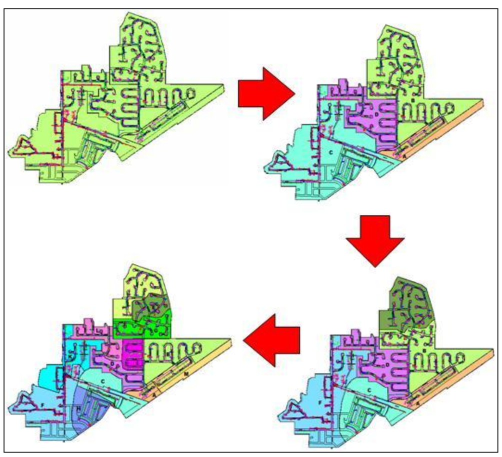

Figure 17 – Fiber Deep Segments a Serving Area

We will recognize the N+0 benefits in subsequent quantization through its effect on aggregate capacity in a serving group, the SNR it enables in fully evolved FTLA architectures, and the new spectrum it tees up for exploitation of new capacity.

Downstream M-QAM Readiness

Using QAM requirements (Figure 4, Table 2) and Table 4 performance for 750 MHz networks with analog loading, we can derive what M-QAM bandwidth efficiency can be delivered to CMs over a range of HFC and home architecture variables prevalent in a typical network deployment. This is shown in Figure 18. Shown are the above calculated CCN values of 42 dB (N+6), 44 dB (N+3) and 47 dB (N+0), labeled via the pink vertical lines. Any other relevant spectrum/loading scenario can of course be evaluated.

Black horizontal threshold lines represent variations of the drop/home architecture, in this case assuming a fixed tap port level of 15 dBmV (analog reference). Different drop lengths have different loss, and the amount of splitting at the home ahead of the CM also varies the level into the receiver. An assumption for architecture evolution is a Point-of-Entry (POE) Gateway-type approach, and we therefore limit the splitting to 4-way maximum. This is a *major* assumption for the evolved home architecture with important implications associated with the CM SNR contribution. For an assumed CM Noise Figure of 10 dB and a Tx/Rx MER (effectively all non-channel implementation losses) of 43 dB, we can observe the range of M-QAM formats that can be supported across the variables shown.

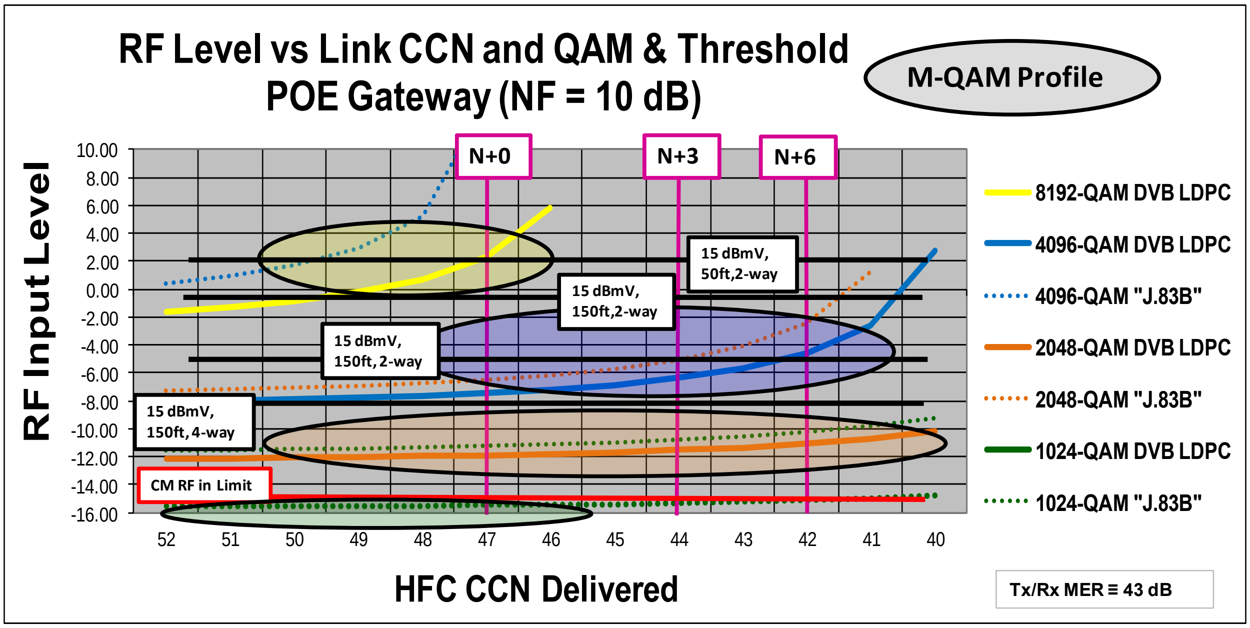

Figure 18 – HFC Geography and Home and Plant Architectures Means a Range of SNR

The QAM + FEC profiles that suit this set of conditions are identified by the colored ovals, which circle the region of operation of the combined variables. In this case, depending on where in the plant a subscriber was located and what drop/home architecture exists, four different QAM formats might be obtainable. Again, we have not discussed margin, but for a simple "margin" philosophy, consider that each CM reports an SNR, and the CMTS selects the next lower (more robust) profile as margin. In such an example, the same number of formats would exist, but instead of ranging from 1024-QAM to 8192-QAM, they would range from 256-QAM to 2048-QAM.

Key items that Figure 18 reveal under these specific assumptions and a relatively limited range of drop/home variables are listed below. Note also that tap port levels – fixed in this analysis – are difficult to keep aligned to a very small range over a series of taps in a string.

1) The plant capability, at least, puts 8192-QAM in play as a possibility. It is not the limiter as architecture evolution continues either to N+0 or to a Remote PHY approach.

2) There is only minor sensitivity to the range of HFC performance for 4096-QAM out to at least N+6. There is almost no sensitivity for 2048-QAM. The drop/home architecture is the more significant factor for the NF assumed.

3) The sensitivities to HFC CCN would increase with lower NF CPE, but the average modulation profile possible for a given drop/home architecture would also be more efficient as a result of the lower NF. An 8 dB NF may be reasonable without excessive CM cost burden.

4) 1024-QAM with a J.83B flavor of FEC would actually be achievable today.

The evolution from 750 MHz architectures to 870 MHz architectures to 1 GHz architectures has by and large been about expanding the bandwidth with new optical and RF technology while achieving equivalent EOL worst case noise and distortion performance. Thus, Figure 18 is relevant, though not exact in the better performance cases, as the forward band extends to 870 MHz and even to 1 GHz. The conclusion – a range of SNR performance – remains, with the range similar but slightly compressed (on the order of 1 dB) due to the minimum performance being the same, and the better performing scenarios – the shorter cascades – degraded from 750 MHz performance by the heavier loading.

A similar statement can be made for RF loading as digital loading replaces analog loading. Clearly discrete distortions such as CSO and CTB reduce considerably in exchange for more digital distortion components, and their contribution would be reflected in CCN or MER degradation that must be managed. Again, variations are small and mostly would be reflected in the better performing cases since end-of-line targets are typically non-negotiable minimums.

Intriguing about Figure 18 is the range – there is obviously capacity left on the table if a 4096-QAM set of users is only receiving at 1024-QAM or 256-QAM. Is there a better approach? We discuss this in the next section.

Something New: Switched Broadcast

Figure 3 and Figure 18 tell us very similar things from different perspectives. In Figure 18, we have the HFC plant telling us the channel quality is there to significantly improve bandwidth efficiency if we can get sufficient level to the home. Left unsaid by Figure 18 is that, assuming we are looking to improve capacity through higher order M-QAM profiles, it would be entirely reasonable to expect CM sensitivity to improve over what is achieved today. The same can be said of fidelity requirements of the equipment on both ends, but, again, Figure 18 does not tell us anything about that. It just says –"I, your HFC plant, can deliver your SNR requirement if you care to exploit it."

Figure 3, however, does include fidelity component as well as the plant component. It just is based on "legacy" fidelity components. And yet, still, it is reporting to us "I have an SNR reserve for a lot of my modems that can do better than 256-QAM."

Recall, in Figure 3, we identified a 2σ range (> 95%) of CMs of about 6 dB. Therefore without even attempting to accommodate the other ~5%, we can identify three QAM thresholds in 2σ – 6dB being one square order apart, and the half-step QAM format between. For example the range 1024-2048-4096 QAM is a 6 dB range. The 4096-QAM users are losing out on some possible capacity if they are only running at the 1024-QAM level.

The above is the basis of a switched broadcast approach to the downstream, and Figure 19 captures this with the use of Figure 3. This is also referred to as Multiple Modulation Profiles [15,18]. As described, Figure 3 is somewhat the CM equivalent of Figure 18, but with the equipment limitations built into the SNR reported. This fact is good in that it is a very practical representation of reality. At the same time, because it is an MER measurement, it may capture equipment effects that specifically are insufficient for needs beyond 256-QAM. And, the home variations that may not exist in a POE deployment are also represented. These differences would skew positively the Figure 3 distribution.

In any event, because we *do* care about the remaining 5% of CMs not in the 2σ range – both above and below, we can, for example, split Figure 3 up as shown in Figure 19. This example creates five regions. More regions can be created to cover the distribution with more granularity, as shown in the black "dashed" lines creating two intervals between the colored ones shown. Such granularity is not easily available with modulation formats without going outside of the rectangular M-QAM family, but can be achieved through the use of different code rates of the LDPC FEC. Again referring to Figure 3, the relationship of lower code rates leading to lower SNR requirements is possible to determine if this granularity is desired.

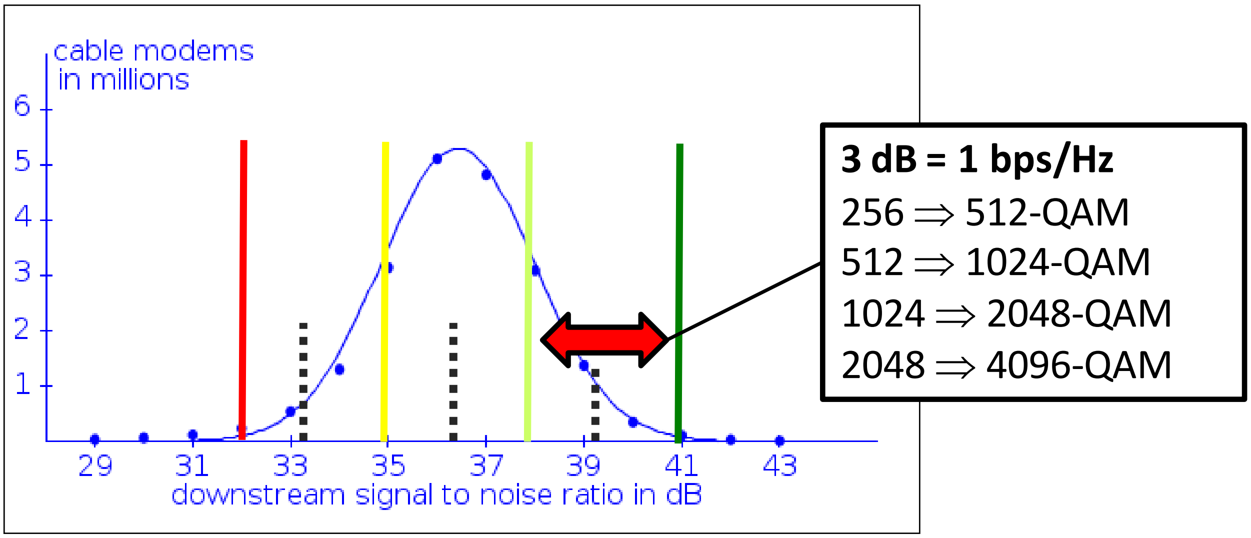

Figure 19 – Multiple Modulation Selections Exploit the Range of CM SNRs

If a CM has a choice of M-QAM profiles, then it can select the one that optimizes its capacity. Every CM can do this, and select among the QAM-FEC "buckets" that suit its estimated channel performance, optimizing the average capacity usage. Each bucket represents a group of modems subscribed to the same "broadcast" profile in what amounts to a switched broadcast downstream.

It is straightforward to make a simple estimate of the relative gain above a single 256-QAM broadcast selection using the above distribution and choosing the average M-QAM profile as 2048-QAM. Based on Figure 3 and Figure 8, the threshold for 2048-QAM is about 32 dB SNR. If we use 35 dB for margin purposes (one QAM profile lower than what a modem reports gets selected), then this threshold is about 1σ lower than the statistical mean of the distribution in Figure 3. The calculation easily follows, with 256-QAM being a vanishingly statistically small percentage but which must be upheld as a fall-back QAM profile, leading to the following:

- ~16% @ 1024-QAM
- 68 % @ 2048-QAM
- 16% @ 4096-QAM

This of course averages 11 bits/symbol, or 11/8 = 37.5% capacity gain. Using a margin of 6 dB, similar for example with what is used today for single-profile 256-QAM, reduces this efficiency gain to about 25%. Note that the maximum efficiency gain is 50% at 4096-QAM – 12 bits/symbol of 4096-QAM to 8 bits/symbol of 256-QAM. Per Figure 18, it appears with N+0 – and probably some enhancements to receiver performance – 8192-QAM is within reach, and up to 62.5% spectral efficiency maximum gain.

The calculation with increased margin identifies one of the other significant advantages of a switched broadcast approach as opposed to traditional broadcast of 256-QAM only. The 256-QAM only (or any single profile selected) must be able to be received by all robustly to be an effective solution. It therefore enforces a "lowest common denominator." Whatever the least capable CM can achieve is what everyone receives. Since there are often outliers and corner cases, these lowest-performing devices drive the total channel capacity as well as the margin allotted to ensure robustness. With no other option to handle a connection problem other than a truck roll, operators tend to ensure a very conservative field margin when a CM is deployed [18] – again, about 35 dB as a typical number for a receiver slicer threshold of 27.5 dB and a DOCSIS minimum BER/SNR requirement set at 30 dB SNR.

By contrast, a switched broadcast enables a CM that is experiencing connection problems – such as counting excessive codeword errors – to switch down to a more robust profile and (likely) continue to have service. This can work in the other direct as well – as plant evolutions occur or home architecture re-engineered, modems can move to more capable profiles for more throughput. As such, because plant variations of performance tend to be very slow under normal conditions, periodic updating of the QAM buckets is possible, and interrupt mechanisms potentially permissible when problems ensue, it is foreseen that the need to run the downstream operation with 6-10 dB of margin for robustness will no longer be necessary. A reasonable recommendation, for example, again might indeed be to choose a profile that is one half-step more robust than the profile that a CM reports that it can support.

Figure 20 diagrams how a switched broadcast approach may operate in practice. The approach does not come for free. There is an obvious increase in complexity in the MAC scheduling function, which now must schedule groups of modems instead of blasting out traffic with little knowledge of the receiving aggregate.

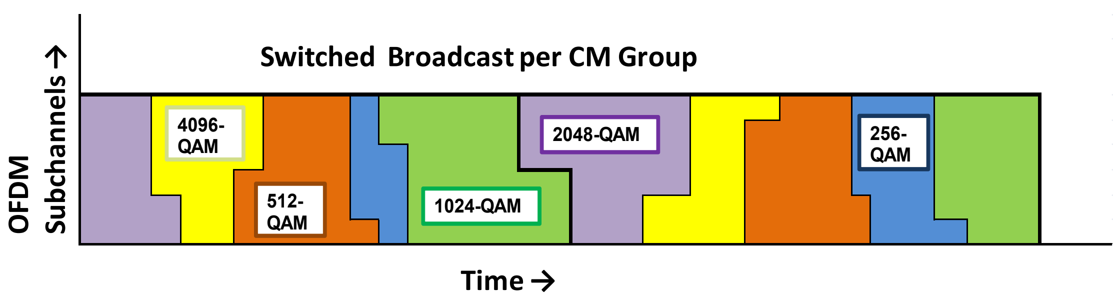

Figure 20 – "Switched" Broadcast Modulation Profiles Exploit the Range of CM SNRs

Additionally, a major element of how efficiently an OFDM channel can be used is the choice of cyclic prefix (CP) and its relationship to the symbol time that was identified previously. In short, the CP manages the reflection energy by basically waiting out multipath. CP is selected to outlast the echo, but in so doing removes time of payload transmission from the channel, costing efficiency.

To manage the complexity of switched broadcast, the same CP for each profile segment should be used. Similar to the SNR least common denominator that drives the idea in the first place, we are now subject to a CP lowest common denominator. The CP does not actually have to be chosen to completely outlast all of the echo content – residual intersymbol interference (ISI) is acceptable so long as it does not significantly affect the total SNR. However, the SNR degradation due to residual ISI of CP must be tolerable for the *highest* QAM profile. Lower QAM profiles may have been capable of a shorter CP and a more efficient usage because they could have tolerated more SNR degradation due to residual ISI. This causes a quantifiable loss of efficiency.

Upstream 85 MHz: Ready and Able

A key component of the evolution of HFC is enhancing the upstream. For many years, it has been recognized that to DFB return optics is required to take advantage of DOCSIS 2.0 and DOCSIS 3.0 capabilities, in particular around 64-QAM at 5.12 Msps. With DFBs assumed coming into place virtually everywhere a high-capacity upstream is desired, and DFB technology having advanced considerably since earlier generation lasers, we can now earnestly look at the ability of today's return optics to support beyond 64-QAM modulation and beyond 42 MHz of spectrum. The limitations of only 37 MHz of upstream, especially with a significant portion of it heavily polluted demands that a wider band upstream be available in the future. DOCSIS 3.1 sets 1 Gbps of upstream as an objective.

In Figure 21, typical performance of an upstream DFBs at nominal link length over the 85 MHz mid-split bandwidth is shown. Also included is the RF noise contribution of a deep cascade (N+6) combined four ways (dashed blue). CMTS receiver sensitivity for high-sensitivity DOCSIS 3.0 upstream receivers is also included to arrive at a net channel response (solid blue).

New PHY performance thresholds using LDPC FEC assumptions (MoCA™ Short) of Figure 7 are shown on Figure 21 with 6 dB of margin to allow for burst receiver implementation complexities and operating margin for the more dynamic upstream channel environment. As in the downstream, system performance suggests we can be much more capable than we are running in most upstreams today.

Of course, the upstream picture is a little more variable than the downstream. There are still major performance limitations where FP lasers still exist and where Headend combining to limit port counts effectively combines noise and halves the available SNR per each combine. Gradually, however, we expect these situations to melt away and be left with DFB links (or digital return, roughly the same performance), uncombined, and over smaller serving groups that also begin to take a bite out of the upstream additive interference problem. Figure 21 allows us to see where this potentially takes us.

It shows that with new LDPC-based FEC, 1024-QAM is possible for high performance DFB optical links available today over the full 85 MHz. There is about 13 dB of dynamic range (DR) above the threshold – not as much range as today's 64-QAM over 42 MHz but above the 10 dB DR standard typically used to define sufficient robustness for the link itself. Clearly, for the given link performance, the FEC is making an important difference for 1024-QAM support.

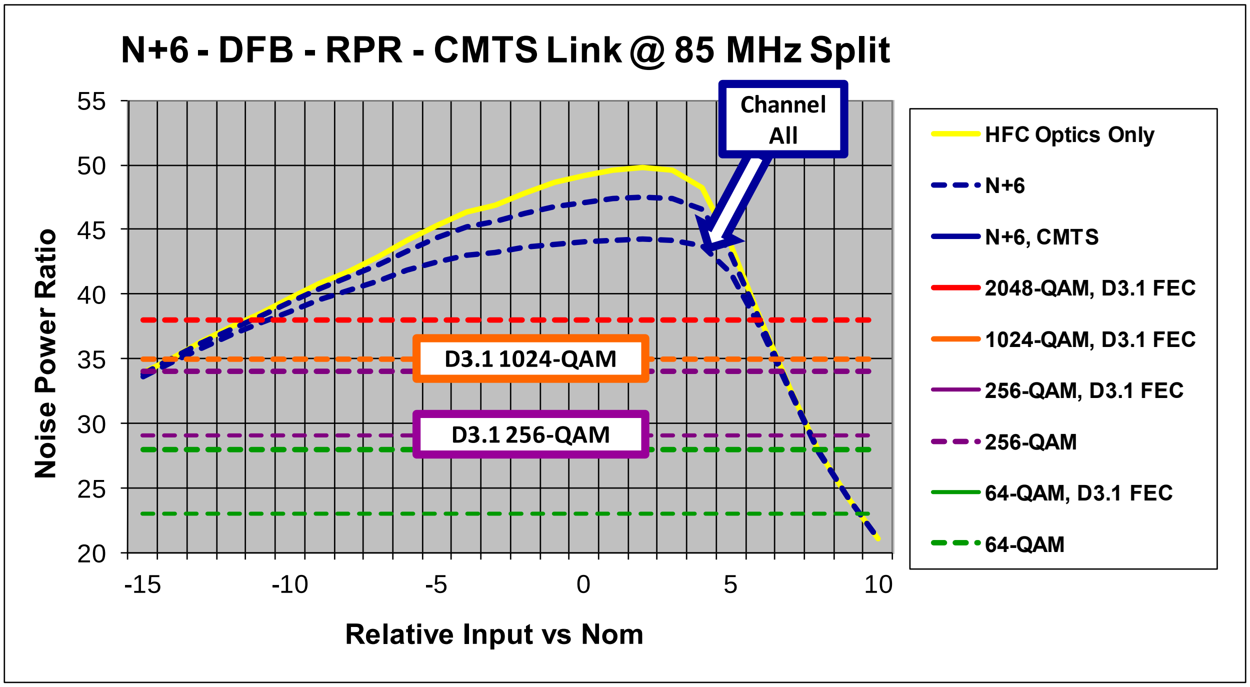

Figure 21 – Modern DFBs, Improved CMTS Sensitivity, 85 MHz of Spectrum, and New FEC Create New Upstream Capacity

As Figure 21 also shows, 2048-QAM has precisely 10 dB of dynamic range, so is actually borderline sufficient. Extended link lengths, minimum guaranteed performance, or older DFB (1 mw) lasers might yield insufficient dynamic range for typical robustness. Also, not all CMTS receivers are created equal, and without a high sensitivity receiver, net end-to-end DR will be degraded. Lastly, the complexity of 2048-QAM itself probably demands an allowance for additional implementation loss and/or dynamic range threshold. It would be very premature to state that the upstream is capable of 2048-QAM until further analysis can be done and bust receiver artifacts better understood. But, that the link is in the ballpark of good enough is encouraging.

Figure 21 indicates why all of the M-QAM formats in Figure 2 are worth considering for the upstream as well as the downstream. The technology and architecture variables are falling into place to make these possible from a plant perspective, shifting the performance burden to the complex task of burst receivers.

Up, Up, UP and Away

High split describes essentially anything that extends beyond the DOCSIS-defined 85 MHz mid-split, but has for many years implied an upstream spectrum of 200 MHz. In Figure 22, we extend the prior analysis to this case of a 200 MHz split. All laser characteristics are assumed the same, so the calculation is based only on signal loading loss associated with the sharing of a fixed power into the laser over a wider bandwidth.

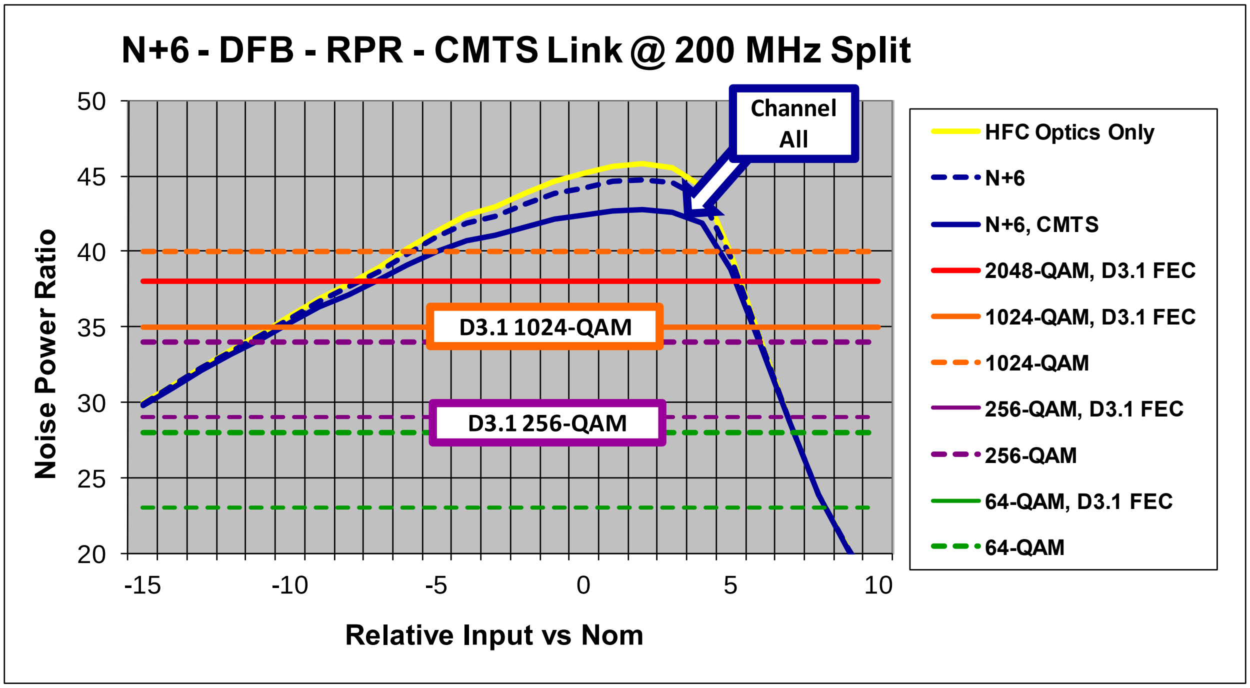

Figure 22 – Extending the Upstream to 200 MHz

Encouragingly, without any improvements assumed aside from the laser itself performing identically over the wider bandwidth, the 256-QAM mode is supported robustly despite the loading loss with plenty of dynamic range. This should not be too surprising since 256-QAM has been proven in the upstream using today's technology [5,6].

The 1024-QAM case still has sufficient DR, but it is now a borderline case at exactly 10 dB under typical conditions using today's DFB technology at nominal length. It would likely not scale in every situation as adequately robust. We can see that 2048-QAM now has insufficient dynamic range as well as a low operating margin of about 3 dB.

On the bright side, there is nothing to suggest we are out of hope expecting to get 1024-QAM (or higher) across a 200 MHz linear return link. We are within single digit dB ranges of achieving key robustness objectives – the kind of dB differences that technology evolution usually overcomes with time and development.

Based on the evolution of technology (high power DFBs), modulation profile (256-QAM), and spectrum (85 MHz proven and 200 MHz projecting well), the upstream is well along the way down the path of achieving a 1 Gbps objective and covering an expanded spectrum range with high bandwidth efficiency. Furthermore, 1024-QAM upstream appears immediately within reach with new FEC and modern DFBs, consistent with the fact that 256-QAM can be achieved today with "old" FEC.

Lastly, robust 2048-QAM from an HFC link performance does not seem like a stretch already, and is in fact borderline acceptable from the plant SNR perspective for the 85 MHz case.

<u>New Capacity = New Spectrum</u>

Equations 1 are 2 were pretty clear about the role spectrum plays in finding new capacity. We know already that the best we can expect from spectral efficiency is about 50% in the downstream with 62.5% perhaps attainable eventually. In the upstream, these numbers are 67% (1024-QAM) or up to 83% (2048-QAM).

Most MSO's concern is currently around downstream because of the persistent aggressive growth. We observed this in Figure 1, and recognized how threatening this could become. Meanwhile, upstream CAGR has stagnated, putting little pressure on the urgency of relieving the inherent spectrum bottleneck of 42 MHz. As segmentation is occurring, driven by the downstream, the benefits of average upstream bandwidth per home are available to the upstream as well, assuming it is simultaneously segmented as is usually the case. However, the benefit does not actually always accrue, because this is often handled by a combining function in the HE until the traffic demands a new upstream port.

The percent new bits per second of capacity and the lifespan they represent can easily be converted to lifespan metrics through the concept of Traffic Doubling Periods (TDPs). Some simplified relationships are shown in Table 5 below.

Table 5- Traffic Doubling Period Relationships

TDP (years)	CAGR %	Simple %
1.7	50	50
2	41	40
3	26	25
4	19	20
5	15	15

Table 5 is very useful for back of the envelope calculations in the range of CAGRs meaningful to cable. Obviously, if the downstream is growing at 50% CAGR and we add 50% more capacity immediately tomorrow, that step is worth about 12 months of lifespan (wow, is that all?). However, if it settles to about 40% CAGR and we add 62.5% capacity, then we add about 17 months (still, is that all?). Indeed, while these do not sound like much in isolation, this is the nature of trying to deal with the exponential (growth) with the linear (bandwidth efficiency enhancement). This is why a set of tools and techniques must be considered. For example, the picture is less ugly as we saw in Figure 1 when segmentation is included in the equation. Each segmentation is equivalent to one TDP.

And, the situation is less ugly in the upstream. Even at 25% (high), a couple of node splits means 3 + 3 + 3 = 9 years of lifespan without any improvement in spectral efficiency or new spectrum. And, there is actually more spectral efficiency gain available percentage wise simply because the QAM profiles begin lower.

We will delve back into lifespan in the next section. One thing quite clear, however, is that spectral efficiency is but one part of the lifespan extension equation, and a relatively modest one at that in some cases. Node splits are incremental business-as-usual methods that deliver more average capacity as well. However, even in this case there are often diminishing returns trying to split serving groups evenly. And, it is well understood that we have hardly used the entire spectrum that can be made available on the coaxial medium. Thus, there is significant interest in finding ways to exploit new spectrum.

Based on this recognition that spectrum is critical to adding capacity in the HFC network, Figure 23 is an example of a likely long term spectrum evolution [1] over time. A possible "final state" for bandwidth allocation on the coaxial cable is also shown, albeit with some ambiguity around the return-forward crossover band. The industry is beginning to settle around a "high-split range in the region of 200-300 MHz.

Note that by using the downstream above 1 GHz, we are extending a relatively well-behaved channel into an uncharacterized area where it will suffer more attenuation, as a minimum. We first saw this in Figure 13. However, the downstream bandwidth may need to increase if only to offset the loss due to growth in the upstream band should it extend to 200 MHz band or greater. This downstream extension is shown in Figure 23 by the block labeled "New NG Forward." The use of OFDM has unique value in addressing this uncharacterized band, as discussed.

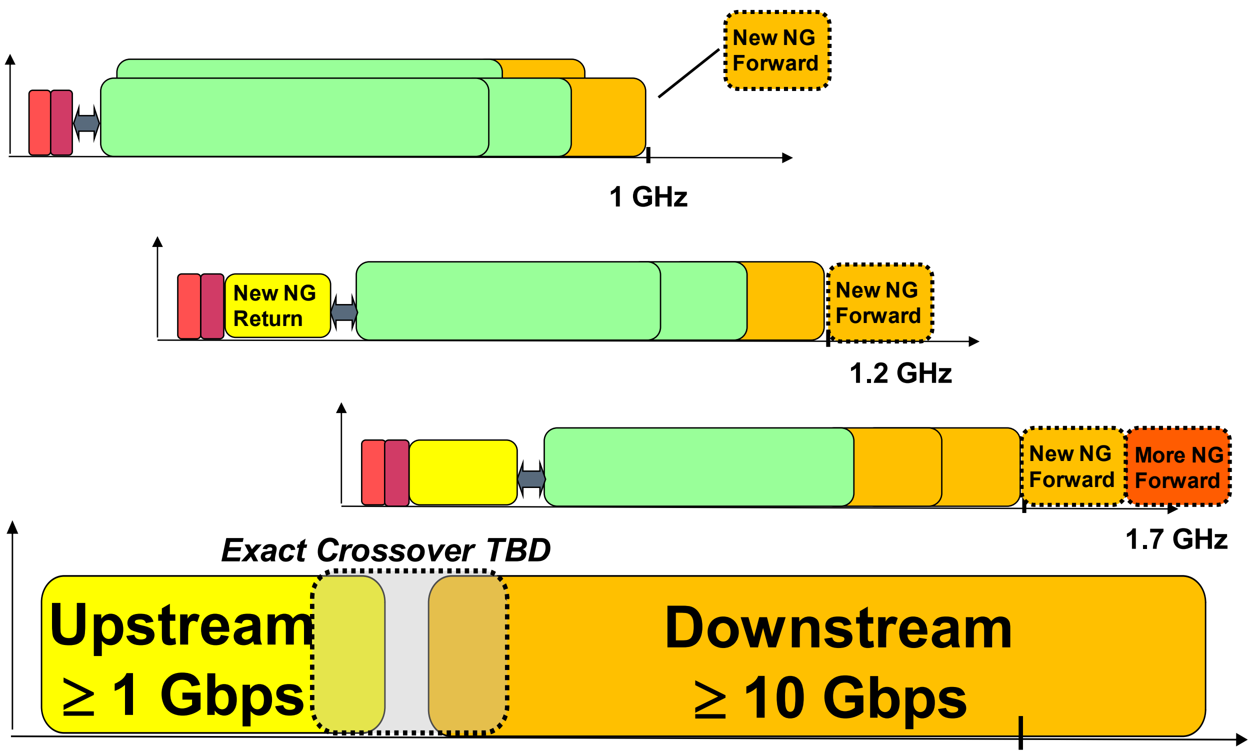

Figure 23 – Possible Long-Term HFC Spectrum Evolution

In the upstream we are instead extending a partially troubled channel into an area where we expect, in general, a better environment. The upstream today gradually becomes well-behaved with increasing frequency above about 15 MHz in North America. As we extend the spectrum above 42 MHz, cleaner bandwidth will become available, enabling more bandwidth-efficient use. The FM band is, of course, an area where characteristics may be less friendly for upstream due to funneling if we extend to 200 MHz. The implications of use of this band must be determined.

Referring to the stages shown in Figure 23, note that the upstream evolution takes place as an extension to mid-split, and subsequently an extension beyond this labeled "New NG Return." The idea is that the 85 MHz mid-split is available in current DOCSIS 3.0 and HFC technology today, and offers a very long window of upstream lifespan and service rate growth to the 100 Mbps threshold. At some point in the architecture migration, the new phase of upstream to achieve 1 Gbps can be introduced. At this point in time, perhaps due to service evolution such as IP Video and legacy removal, the downstream may be prepared to accommodate a loss in spectrum to upstream use. Otherwise, this may be the point in time to extend downstream. This appears to be the more likely scenario, due to the likely slow withdrawal of legacy services and the need therefore to simulcast, burdening downstream spectrum. Initially, the extension may simply be excess bandwidth above 1 GHz such as 1.2 GHz before evolving to a broader chunk of bandwidth exploitation above 1 GHz if necessary.

An example of the "excess bandwidth" of a single 1 GHz tap such as those very commonly deployed today is shown in Figure 24 [3].

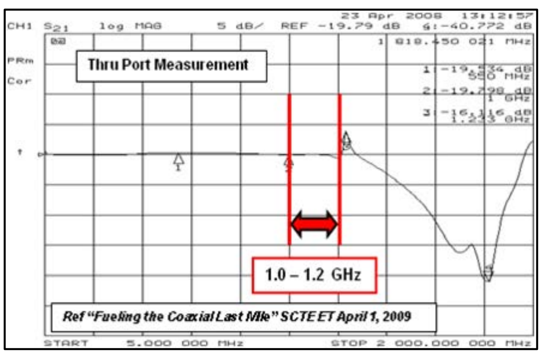

Figure 24 – 1 GHz Tap "Excess" Bandwidth

Note that Figure 24 captures one single tap. In an actual RF leg, there will be multiple taps, and in an HFC cascade, there will be amplifiers and taps following a fiber optic node, and these will all cascade to create an aggregate frequency response. This is important to understand, since it is *not* the case that most actives in the field are 1 GHz. See [3] for further discussion.

Going beyond 1.2 GHz will most likely be necessary to achieve 10 Gbps of useable capacity, in particular with an extended upstream band to 200 MHz. Figure 24 makes it clear that there is not much hope for this on 1 GHz taps, especially in a cascade. taps with wider bandwidth capability are certainly possible. However, removing old taps and replacing them with new ones is time consuming, costly, and intrusive. With a faceplate change option, the ability to convert a 1 GHz tap to a 1.7 GHz tap can occur with minimal down time, decreasing the expense to the operator. Figure 25 shows an example frequency sweep of tap faceplate technology.

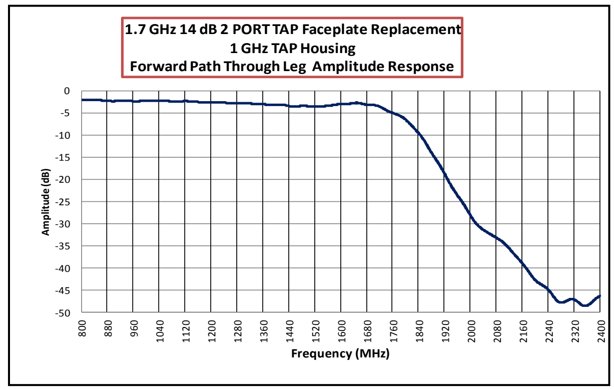

Figure 25 – Creating New Forward Bandwidth with Tap Modification

How long will it take to migrate the HFCs spectrum? The period over which the spectrum evolution in Figure 23 occurs is probably on the order of 15 years. However, there is probably an equally strong likelihood that it is either 7-8 years or never. The play out of service trends and growth will have a major impact on the outcome because of the very large long-term implications of, for example CAGR remaining at 50% or settling to, say 30% and how it is trending at the time. It's the difference between 30 Gbps and 7 Gbps of aggregate capacity after ten years. It's the difference of whether spectral efficiency and service group segmentation is sufficient, or new spectrum is clearly necessary. We will discuss further using the Capacity Management Timeline in the next section.

No Free Launch

While new RF spectrum to exploit is exciting from a capacity standpoint, it has some challenges in practice, since the band where we are adding spectrum is of higher RF attenuation. There is an optical loading component as well, but since transmitter loading is flat, this added spectrum costs only about 0.8 dB relative to 1 GHz performance. Of course, 0.8 dB can be meaningful. As Figure 18 or Figure 19 indicates, for example, the shift has an impact on the percentage of users that fall into a particular M-QAM profile bucket, with the shift skewing modem distributions down and lowering the net capacity increase. It is less than 0.8 dB if we assume that the upstream extends to 85 MHz for example, but the difference is small.

However, the RF spectrum placed on the coax is uptilted for frequency dependent attenuation. Adding new RF load to the high end increases the total power disproportionately. This is easily quantified. Figures 26-28 demonstrates the issue. In Figure 26, a 550 MHz downstream of analog carriers with digital loading to 1 GHz at 6 dB digital de-rate and typical uptilt is shown. The total RF power out of the active with this loading is 70 dBmV.

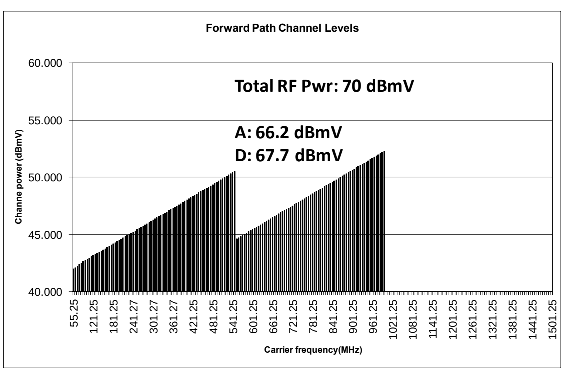

Figure 26 – Typical Downstream Analog + Digital Loading

In Figure 27, the extra 200 MHz of bandwidth is loaded, and just this 200 MHz increases the total RF power load by almost double, at 2.7 dB. Fortunately, GaAs RF technology has given way to GaN technology, which is capable of higher output at equivalent distortion performance compared to typical fielded GaAs amplifiers. This has generally been engineered as higher in-band RF levels across the band by 2-3 dB or used as bandwidth extensions of older plant that maintains active spacings, but it is expected that this technology, now on its second generation, will enable 1.2 GHz at equivalent 1 GHz performance.

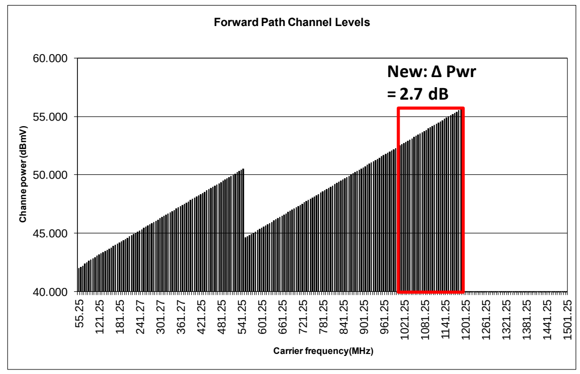

Figure 27 – Tilted RF Outputs Increase Total Power Disproportionately

Now consider Figure 28, where the bandwidth is extended to 1.6 GHz. In this case, implementation of the band by simply loading the spectrum at the identical tilt and at the equivalent relative PSD is not sustainable, requiring 7.2 dB more RF power, over a 5x increase. The 5x interpretation is useful for considering the impact to the DC power budget needed to drive the additional RF power! Note that this 7.2 dB is the case even under the low end loading assumption shown in Figure 28, which is a high split return band with a forward path that begins at around 250 MHz. Also, the optical loading for this case, should laser transmitters achieve equivalent performance parameters is now 2.1 dB – much more significant.

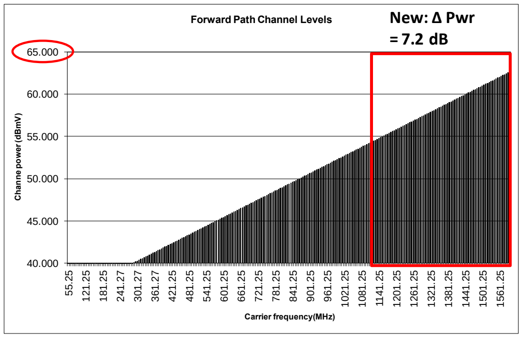

Figure 28 – Downstream Bandwidth Extension to 1.6 GHz Creates an RF Power Dilemma

In summary, the addition of 1.2 GHz of bandwidth looks to be a manageable extension using current techniques for optical loading (or Remote PHY) and RF distribution, including tilted RF outputs used as they are today. This is not the case at this point for enabling up to 1.6 GHz. For this case, it appears use of the band may best be served by the linear optical band extending, for example, to 1.2 GHz as part of the spectrum evolution shown in Figure 23.

Then, at some point in the future should the additional bandwidth be necessary, with the tap capability seeded into the field and the architecture migrated to N+0 so that RF actives are not an obstacle, an overlay solution that implements this additional bandwidth can be deployed, perhaps as part of a Remote PHY strategy. This region is labeled “More NG Forward” in Figure 23, and as a separate overlay would also benefit from the fact that it represents a standalone narrowband solution

ADDING IT ALL UP

Taking inventory of the itemized list of evolutionary techniques that offer the opportunity to breathe new lifespan into the HFC network, we have covered the following:

- Figure 1 – Analog spectrum and eventually other legacy must be assumed removed.
- Figure 2 – Employing advanced M-QAM formats will offer more spectra efficiency and more capacity over a given spectrum.
- Figure 4 – New LDPC FEC allows these more bandwidth efficient M-QAM formats to be possible at SNRs that exist over today’s HFC network performance as shown in Figure 8.
- Figure 9, Table 3 – OFDM does not itself provide “new” capacity, but its channel optimizing nature enables the full capacity of M-QAM and FEC to be obtained in all environments and its use allows bandwidth efficient use over an expanded RF spectrum range.
- Table 4 – HFC fiber deep evolution provides increasingly better channel SNR performance for the aforementioned M-QAM and FEC to exploit.
- Figure 17 – Fiber deepest (N+0) maximizes the channel SNR delivered over HFC *and* opens up the opportunity to expand spectrum above 1 GHz for fresh, new capacity. It service group shrinking benefits produce over 90% aggregated BW savings.
- Figure 18 – Analysis of fiber deep architectures combined with home architecture evolution via a POE gateway approach shows 1024-QAM through 8192-QAM within reach, with 8192-QAM as a potential maximum bandwidth efficient profile – a 62.5% increase over 256-QAM in spectral efficiency.
- Figure 19-20 – Switched broadcast allows CMs to be classified into buckets based on their channel quality. This increases average capacity, maximizes a user’s experience, and allows the operator to

eliminate the waste of dB dedicated to margin associated with a lowest common denominator single broadcast profile approach.

- Figure 21-22 – Modern DFBs, high sensitivity DOCSIS 3.0 receivers, shorter cascades, elimination of HE combining, advanced QAM profiles (Figure 2) and new FEC (Figure 4) do for the upstream what similar M-QAM, FEC, and architecture evolution do for the downstream. 64-QAM to 256/1024/2048-QAM (33-83%) is realistic.
- Figure 23-25 – Spectrum evolution frees up room for both downstream and upstream growth. Downstream 1.2 GHz or 1.6 GHz from 950 MHz or 500 MHz total is a 20%-90% increase. The upstream spectrum increase is over 100% (85 MHz) to over 300% (200 MHz).

We can lay these evolutions out on a downstream Capacity Management Timeline as in Figure 1. We can perform a similar Capacity Management analysis for the upstream. We will actually start with the upstream, since a key conclusion from that translate into details of the downstream Capacity Management Timeline calculations. And, as indicated, it is the less urgent of the two at this point for most operators and is generally managed effectively in concert with downstream segmentation.

Lifespan Management – Upstream

Similar to the calculations in the downstream, we can project how long our upstream will last under various growth and service scenarios. We base the projections on 80 Mbps deployed today, such as two 64-QAMs and one 16-QAM at 6.4 MHz.

Unlike the downstream, CAGR for the upstream tends to be less predictable – spiking when Napster and YouTube were introduced for example, and lagging for several years during other periods of time, including the last several years. Because of this and because of the hard limit of upstream spectrum with fewer tools to manage it than the downstream, we can create a form of the analysis that is simpler to interpret. This is shown in Figure 29. With a typically more variable CAGR range, it also handily provides a sensitivity analysis perspective around incorrect CAGR projections.

Figure 29 reveals why both an expansion beyond 42 MHz is not as urgent as it might seem it should be. The recent CAGR range is identified, and it has been relatively modest. Referring to Table 5, even the 25% CAGR has a three year doubling period, so with two node splits along the way in concert with the downstream, six years of latent growth lifespan is available even if fully congested today. This becomes 8 years for the 20% CAGR. This is readily apparent from the nearly full 100 Mbps ATDMA solid yellow trajectory, and the 9-year lifespan after two splits that exists for it. The timing of course becomes the key to those splits. Figure 29 results are best synchronized with Figure 1 for the purpose of the timing of the node segmentations.

Using SCDMA and two splits and a 20% CAGR, more than a 10-yr window is available, with the first split lasting out to 7 years.

With an 85 MHz mid-split, DOCSIS 3.0 technology with a 20% CAGR provides 7 years of life, and with two splits this extends to 15 years! It is for this reason, and that this bandwidth can support a service rate of 100 Mbps – not practical over 5-42 MHz return – that an 85 MHz is potentially a long-term upstream bandwidth phase, and not an “interim” one.

Lastly, with a 1024-QAM capable upstream (red) even a 25% CAGR has 15 years of aggregate capacity life with an 85 MHz mid-split. There is almost a full decade using an aggressive 40% CAGR upstream. Again, this is pretty simple to get a feel for from Table 5: 40% is approximately a 2-yr TDP and 85 MHz will be on the order of 550 Mbps of capacity. 80 Mbps will not quite be three TDPs of 2 years – or slightly less than 6 years. Two splits are two TDPs or 4 years, so slightly less than 6 + 4 or slightly less than 10 years total.

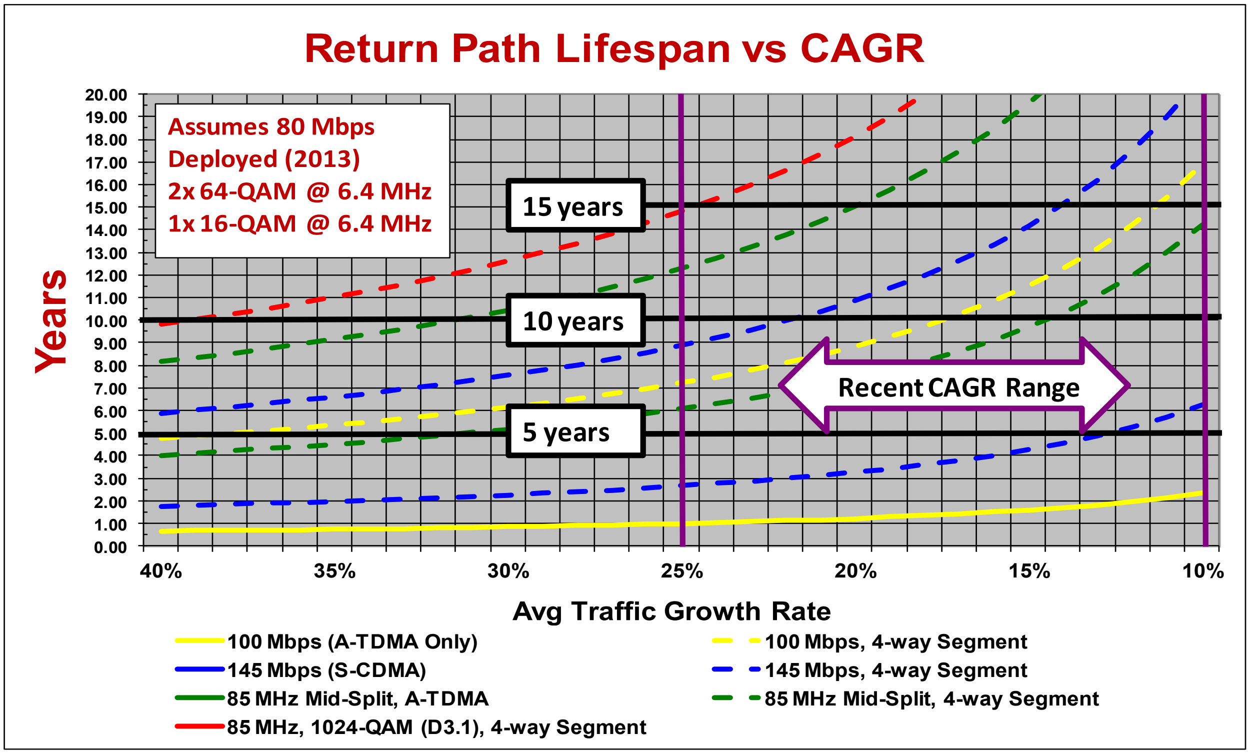

Figure 29 – Upstream Capacity Management

At nominal CAGRs of today, should that continue for many years, a capacity-enhanced 85 MHz upstream offers so much room for new growth in terms of lifespan that it does not make very much sense to plan what the next step should be after it. It does of course make sense to keep a watchful eye on nominal CAGRs. Unfortunately, the 85 MHz mid-split cannot achieve the 1 Gbps DOCSIS 3.1 objective. However, because of legacy constraints around STB signaling, and the squeeze that extended upstream spectrum places on the downstream, it may be unwise to consider a broader extension for some time in any case.

To put 1 Gbps in a CAGR perspective, it is four TDPs to exceed the threshold from 80 Mbps. At 4 years per TDP at today's CAGRs, this means 16 years to the need date from an average consumption standpoint.

A path that first evolves to 85 MHz, covering well over a decade of new capacity, delivering 100 Mbps speeds, and maintaining legacy out-of-band signaling is the most prudent phased approach. In so doing, technology to later enable a very simple switch to a wider band upstream should be deployed to minimize a subsequent costly plant touch, albeit the need for this capacity in the upstream for residential services looks well into the next decade, and many of the actives that include such technology may be removed by fiber deep evolution by the time of need.

Downstream Lifespan: Worth the Price?

An updated version of the Figure 1 downstream Capacity Management Timeline is shown in Figure 30. Because of what we have learned examining the upstream, the spectrum used for forward band allocation is an 85 MHz mid-split, and therefore a 108 MHz forward band starting point. Let's examine what we have achieved between Figure 1 and Figure 30. Note that most of the Figure 1 data is shown on Figure 30, but we have simply extended the timeline for another decade in recognition of, and to evaluate, the additional lifespan we have enabled. To reduce clutter, we only show one of the "All-QAM" thresholds of 256 QAM, and chose the 870 MHz middle case.

Today's downstream spectrum is basically full. Just ask any MSO. Technically, there is no "room for growth." Managing from his completely full state through the balancing act of old services and technology to new services and technology is the challenge addressed in detail in [4,10]. As mentioned in the first section, the objective of this paper was to see how far north we could shift capacity thresholds on the chart given the list of the architecture and technology variables that we are throwing at the problem, and what that buys us.

Using 870 MHz as an example we can make the following assessment based on the blue "IP Video" trajectory (little difference from red long term) and the green "Max Consumption" trajectories, comparing the lifespan ranges. This is shown in Table 6.

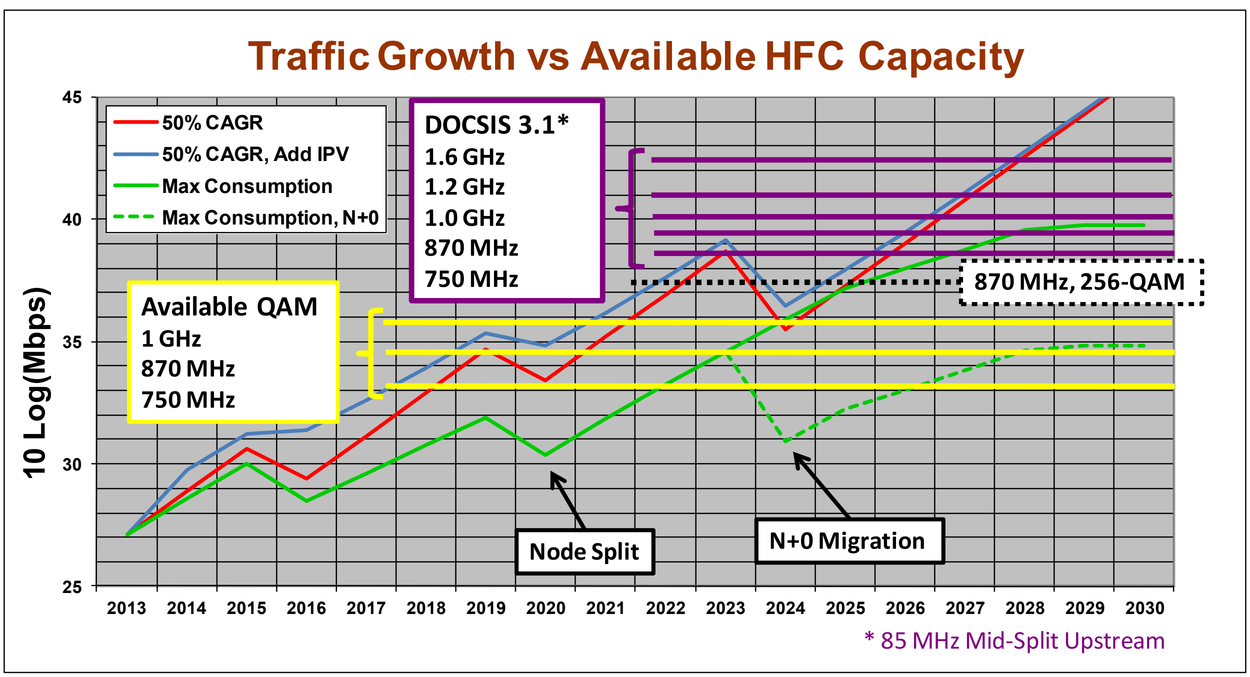

Figure 30 – Capacity Management Timeline: Breathing New Lifespan into the Network

Table 6 – Capacity Management Timeline: Years of HFC Lifespan Starting with 870 MHz

	"Available" 870 MHz	"All QAM" 870 MHz	"New QAM" 870 MHz	"New QAM" 1200 MHz	"New QAM" 1600 MHz
50% CAGR w/IPV					
Two Splits	7	9	10	11.5	13.5
N+0	10	12	13	14	15
Settled CAGR					
Two Splits	10	12.5	14	Sufficient	Sufficient
N+0	14.5	Sufficient	Sufficient	Sufficient	Sufficient

The lifespan increases shown in Table 6 are significant, though not necessarily entirely comforting. Focusing on the top row of Table 6, from the "today" case of "Available QAM," the increases are as small as 3 years with a maximum of 8 years result for the 50% CAGR case. There are a pretty significant number of possible expensive steps invested in to achieve those 8 years. It does not feel like an enormous bang for the buck. The differences are even smaller when compared to the Gbps available in a full spectrum of 256-QAM – "All QAM" – at 4.5 years maximum

This really just comes down to the implications of the relentless nature of the long term aggressive CAGR assumption, if such assumptions turn out to hold. Linear gains to combat exponential processes make it difficult to keep pace. We will consider the alternative in the next section. However, if we convert the absolute years of Table 6 to lifespan percentage gained, at least the perspective looks, perhaps, more in line with expectations.

In Table 7, we tabulate this percent lifespan gain for the three cases of advanced QAM efficiencies for the three spectrum cases relative to "Available QAM" and relative to the capacity of an 870 MHz forward band full of 256-QAM carriers ("All QAM").

Table 7 – Capacity Management: Percent Lifespan Extension, Persistent CAGR

	"New QAM" 870 MHz	"New QAM" 1200 MHz	"New QAM" 1600 MHz
"Available" QAM			
Two Splits	43%	64%	93%
N+0	30%	40%	50%
All QAM			
Two Splits	11%	28%	50%
N+0	8%	17%	25%

For the persistent CAGR case, the combination of tools provides percent lifespan gain that in some of the cases is compelling – from 30% but to as high as over 90% – using what is defined as "Available QAM" as a starting point. This term essentially basically boils down to executing on all the following bandwidth efficiency tools:

1) Analog reclamation
2) New Spectrum
3) New Modulation Profiles
4) Service Group Segmentation

The "All QAM" percent lifespan extension is much more modest. This basically represents the above steps, but not including the analog reclamation as part of the advantage gained. This should not be surprising based on the prior discussion on TDPs and CAGRs. If we can manage 50% new bandwidth efficiency gain, but we are unable to add any actual new bandwidth, than a 50% growth in traffic in a year will consume that gain immediately. This is basically what the one year gained in the columns for "All QAM" 870 MHz and "New QAM" 870 MHz tell us. It reemphasize that new capacity for new lifespan relies on *a set of evolutions* for new technology to combat exponentially growing traffic. Part of this is absolutely the critical removal of existing less efficient legacy technologies, and in particular much less efficient analog services.

In absolute terms, we can see that the most "Northbound" thresholds of 1.2 GHz and 1.6 GHz carry us into the middle of the next decade with two node splits, and late into the next decade for N+0. While the years added in Table 6 do not seem particularly large in most cases, this is an extremely long window of time from a technology standpoint – too far out to do much in the way of strategic planning for capacity management. It is actually a very beneficial observation window for trends in service growth that impact subsequent migration decisions.

Table 6 and Figure 30 prove the obvious – "CAGR forever" always wins the bandwidth battle – it is just a matter of the time scale chosen to find defeat. This is of course a key reason why attention to service growth trends during the window of the next decade or so will be critically important to pondering future architectural evolution. In the case of continued persistence, the only step left likely to be of sufficient ROI is FTTH.

Asymptotic Growth

Perhaps the more intriguing Capacity Management Timeline result and conclusions they imply are the cases of "Max Consumption." These trajectories, drawn from analysis in [7], are based on the premise that streaming video as the driver of CAGR arrives at a quantifiable maximum video bit rate and demographic-based concurrency per-household peak. If this is so, then the CAGR associated with streaming video will taper the traffic growth towards this maximum over the course of time that these video evolutions take place. The green trajectories demonstrate this hypothesis, which effectively suggests that 50% CAGR forever is not practical to sustain, with other reasoning for this settling also described in the first section.

The conclusions for these possibilities are intriguing. In the case of N+0, there will be sufficient bandwidth as long as the network is extended to the 1 GHz point – a standard HFC solution available today. And, this is the case with simply 256-QAM available!

If instead only two splits are used, and no extension to N+0, then we coincidentally arrive at the same answer of 1 GHz minimum of spectrum, but in this case we must take advantage of increased bandwidth efficient modulation formats discussed herein.

Importantly, though, the extension to either 1.2 GHz or 1.6 GHz is not required under this assumption of asymptotic CAGR, at least not for the end state – but perhaps for managing the "simulcast bubble." These cases of sufficient bandwidth are labeled as such in Table 6.

Moreover, when the settled CAGR cases (about 35% average is used) *do* represent enough growth to breach a threshold of capacity, they are nonetheless extending the lifespan by a meaningful, if not huge, 3-4 years, and this is with just a relatively modest CAGR setting. This is evidence that the power of compounding, while potentially very threatening, can also be quite forgiving if its most aggressive tendencies settle.

SUMMARY

The persistence of compounding data growth demands that the industry respond in order that its service delivery capabilities are not compromised. Several avenues of evolution are available, and all may need to be deployed in order for the network to be a sustainable service delivery platform for the long term. In composite, they put cable on a path do match or exceed PON targets capacity and bandwidth per user. We have observed that only business-as-usual plant segmentation, part of any logical plan, may buy time but appears insufficient to ensure a comforting lifespan for MSO business planners based on growth trends today. We introduced the Capacity Management Timeline as a way to visualize and plan service growth and architecture investment.

We have looked at each of the key avenues available for capacity growth in detail. To increase capacity, the answers are simple as Shannon has shown in (1). We need more bandwidth and higher SNR.

With higher SNR, we can use more spectrally efficient modulation profiles. To most aggressively exploit these profiles, we need Forward Error Correction that maximizes the profile we can use for a given SNR. To maximize SNR, we need architectural evolutions that deliver the highest performance from the HFC channel, which means taking the fiber as deep as possible – and such that we advantageously also share it over fewer subscribers.

Also, we need more bandwidth to operate these high efficiency M-QAM profiles over. There are two components to more bandwidth – removing the legacy services and finding new spectrum.

In the paper, we described how all of these technologies and evolutions mesh with the cable network and quantifiably described the possibilities they create for new capacity.

Armed with these tools, we put together a longer term picture of the Capacity Management Timeline. We determined that the upstream appears currently less threatening, and quantified the lifespan possibilities for the fully evolved downstream. We observed the extension of lifespan we might expect under a persistently aggressive CAGR with our new capacity possibilities, and learned the obvious – relentless CAGR always wins, only the time of defeat changes.

However, we also learned under specific assumptions that we can add up to 8 years on the lifespan – a long observation window to see just how persistent this CAGR beast is. And, perhaps most interestingly, we learned that with some modest CAGR settling, the HFC lifespan outlook clears up nicely, and offered reasoning why this could be the true future of cable service delivery.

Furthermore, we learned that with settled CAGR and N+0 evolution, "HFC forever" emerges. And, even the two split case does not necessarily require spectrum expansion in the long term. An end state that survives indefinitely appears achievable, but we simply cannot be certain until we have recorded more years of service evolution and CAGR history. If so, then perhaps the greatest challenge will have turned out to be the challenge we face today – surviving the simulcast bubble by managing service and traffic growth on top of old services that will prove difficult and complex to retire.

ACKNOWLEDGEMENTS

The author would like to thank the following individuals for their contributions to the content of this paper, providing invaluable technology and industry insights:

Rob Thompson
Jack Moran
Chris Jeblonski
Phil Miguelez
John Ulm
Fred Slowik
Mike Emmendorfer
John Chapman (Cisco)
Dr. Tom Kolze (Broadcom)
David Urban (Comcast)
John Chapman (Cisco)
Matt Schmidt (CableLabs)

REFERENCES

[1] Chapman, John, Mike Emmendorfer, Dr. Robert Howald, Shaul Shulman, *Mission Is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS*, 2012 Cable Show, Boston, MA, May 23-25.

[2] Howald, Dr. Robert L, *Boundaries of Consumption for the Infinite Content World*, SCTE Cable-Tec Expo, New Orleans, LA, October 20-22, 2010.

[3] Howald, Dr. Robert L, *Fueling the Coaxial Last Mile*, SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[4] Howald, Dr. Robert L, *Looking to the Future: Service Growth, HFC Capacity, and Network Migration*, 2011 Cable-Tec Expo Capacity Management Seminar, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Atlanta, Ga, November 14, 2011.

[5] Howald, Dr. Robert, *The Grown-Up Potential of a Teenage PHY,* 2012 Cable Show Spring Technical Forum, Boston, MA, May 21-23.

[6] Howald, Dr. Robert L, *Time on Our Side: Using the Past to Pave the Way Forward*, The Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[7] Howald, Dr. Robert L, Dr. Sean McCarthy, *Accounting for Techies: Taking it to the Ultra*, The Cable Show Spring Technical Forum, June 10-12, Washington, DC.

[8] Howald, Dr. Robert L, Dr. Sean McCarthy, *Bits, Big Screens and Biology*, The Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[9] Howald, Dr. Robert**,** Michael Aviles, and Amarildo Vieira, *New Megabits, Same Megahertz: Plant Evolution Dividends*, 2009 Cable Show, Washington, DC, March 30-April 1.

[10] Howald, Dr. Robert L. and John Ulm, *Delivering Media Mania: HFC Evolution Planning,* 2012 SCTE Canadian Summitt, March 27-28, Toronto, ON, Canada.

[11] Ulm, John and Gerry White, *Architecture & Migration Strategies for Multi-screen IP Video Delivery***,** 2012 SCTE Canadian Summit, March 27-28, Toronto, CA.

[12] http://maninair.in/orthogonal-frequency-division-multiple-access-ofdma/

[13] http://www.dipol.ie/newsletter2/inf_dipo_2009_32.html

[14] http://www.ict-redesign.eu/fileadmin/documents/ReDeSign-D08_Report_on_Cable_Channel_Model.pdf

[15] http://www.ieee802.org/3/bn/public/adhoc_mmp/adhoc_mmp_urban_01.pdf

[16] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf

[17] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_02_0313.pdf

[18] http://www.ieee802.org/3/bn/public/nov12/schmitt_01a_1112.pdf

[19] http://www.lightwaveonline.com/articles/print/volume-26/issue-9/industry/ofdm-to-power-high-bit-rates-in-next-gen-optical-access-networks.html.

[20] http://www.ni.com/cms/images/devzone/tut/a/bdc8dcae1315.gif

APPENDIX –UPDATED HFC CHANNEL MODEL [16]

The IEEE 802.3bn Working Group has adopted an HFC Baseline Channel Model for purposes of developing a system design for Ethernet PON over Coax (EPoC) technology. This channel model is included here for reference. While cable channel models have been developed in the past, a primary focus for this Baseline was to arrive at "Typical" values for channel parameters, rather than exclusively limit or worst-case parameters. While the limit cases are also part of the effort, it was deemed very important for the task at hand – how much capacity can be squeezed out of the HFC network – to better understand the network's ability to "scale up" in capacity.

The table below is to allow simulation teams to develop and optimize system designs suited to maximizing link capacity. Other scenarios are shown in [17].

Notes	
1	If not defined otherwise, assume typically behaving link but where the behavior is the worst (freq, location)
2	Frequency dependence of coax for broadband calculations: Loss B (dB) = Loss A (dB) x SQRT(B/A)
3	Reference virtual port level for 6 MHz signal at 1 GHz; 15 dBmV Tap port level, 100 ft drop, 2-way splitter
4	(Max Freq - OFDM BW) spectrum range used for drop loss
5	Small drop slope effect on calculation
6	SCN includes HFC geography impact (location in cascade depth)
7	50 kHz Subchannel Reference, Live Video, fully contained within subchannel
	Subcarriers with Interference (50 kHz): Every 70 subcarriers, a cluster of three interfers: I_0, I_0+25 kHz, I_0-25 kHz
8	Typ = CTB/CSO Worst Case Freq; Good CTB/CSO in low-distortion band, Analog contiguous at low end of band
	NCTA measurement method (avg); Error rate simulation should account for PAR and peak durations
9	Worst spectrum regions for CTB and CSO are not the same
10	D/S Burst Characterization in process; BW based on percentage of errored carriers in 8-Channel wide DOCSIS CM
	Duration based on large scale CM sweep of UCER with known interleaver settings; Levels per ReDesign channel model
11	Laser Clipping PSD captured in SCN for out-of-band EPoC Signals
12	Typical tilt, first tap, not equalized, 50 ft drop assumed (Minimum drop impact)
13	Echo mask range for a Single Dominant echo - Does not imply an assumptions about multiple echoes.
14	Meas@700-800 MHz, representive of 99% of modems

System Description				
HFC D/S Spectrum		1.0 GHz		
Cascade Depth		N+3		
Channel Loading		48 Analog (32 removed for D3.1) + 75 Digital		
Optical Architecture		Linear Optics 1310 nm (nominal link length) Legacy		
Home Architecture		Up to max drop length & 2-way splitter		
	#	**Parameters**	**Baseline Channel**	**Notes/Dependency**
Spectrum	1	**Frequency range**	54 MHz - 1 GHz	Note 1
	2	**OFDM Bandwidth**	192 MHz	
RF Level	3	**OFDM Power at CPE Input (dBmV)**		Notes 2-4
		6 MHz BW	-2	
		24 MHz BW	4	
		96 MHz BW	10	
		192 MHz BW	14	Note 5
SNR	4	**SCN Ratio (Signal to Composite Noise Ratio)**	44	Note 6
		Variation over 6 MHz BW (dB)	N/A	Reference Basis 6 MHz
		Variation over 24 MHz BW (dB)	1.5	
		Variation over 96 MHz BW (dB)	2.5	
		Variation over 192 MHz BW (dB)	3.0	
Interference				
Narrowband	5	**CTB Interference (20 kHz BW)**		Notes 7, 8
		# of interfered subcarriers @ 35-40 dBc	0%	
		40-45	1%	
		>45	0%	
	6	**CSO Interference (20 kHz BW)**		Note 9
		# of interfered subcarriers @ 35-40 dBc	0%	
		40-45	0%	
		45-50	2%	
		>50	0%	
	7	**Narrowband Interference (Other)**		
		Bandwidth (MHz)	N/A	
		Level, dBc (PSD)	N/A	
Wideband	8	**Burst Interference**		Note 10
		Bandwidth (MHz)	30	
		Level, dBc (PSD)	-20	
		Duration (usec)	16	
		Period (Hz)	Infrequent	
	9	**Impulse (white) Noise**		Note 11
		Level, dBc (PSD)	N/A	
		Duration (nsec)	N/A	
		Period (kHz)	N/A	
Freq Response				
Amplitude	10	**Amplitude Slope**		Note 12
		dB/MHz	0.01	
	11	**Amplitude Variation**		SCTE Definition, Echo not included
		(dB pk-pk/6 MHz)	1	
		(dB pk-pk/24 MHz)	3	
		(dB pk-pk/192 MHz)	5	
		(dB pk-pk/Total DS BW)	9	
Phase	12	**Group Delay Variation, nsec**		
		Over 24 MHz		
		Mid Band	25	
		Band Edge (24 MHz)	145	
		Over 192 MHz		
		Mid Band	200	
		Band Edge (24 MHz)	320	
Echo	13	**Echo Profile, dBc**		Notes 13, 14
		.5 usec	-20	
		1 usec	-25	
		1.5 usec	-30	
		2 usec	-35	
		3 usec	-40	
		4.5 usec	-45	
		5 usec	-50	
Spurious Modulation	14	**AM/Carrier hum modulation %**	3%	

Upstream

Note this Table, by choice of priorities of interest, represents a remote demodulation architecture (no linear optical return).

System Description				
HFC U/S Spectrum	300 MHz			
Node Architecture	N+3			
Channel Loading	Remote Tx/Rx			
HE Architecture	N/A - EPON Return			
Premise Architecture	Two Way Combining			
	#	**Parameters**	**Baseline Channel Conditions**	**Notes/Dependencies**
Spectrum	1	**OFDM Bandwidth**	192 MHz	
	2	**Frequency range**	100-292 MHz	
Path Loss	3	**Path Loss (dB)**	44	Max loss to first active
		Variation Freq, 24 MHz BW	1	Note 1
		Variation Freq, 96 MHz BW	2.5	
		Variation Freq, 192 MHz BW	5	
Added Noise	4	**Input Noise PSD**	- 115 dBmV/Hz	Contibutions of
Interference	5	**FM Band Interference**		
Narrowband		Bandwidth	8	Spectrum Overlap
		Level, dBc (PSD)	-40	Note 2
	6	**Common Path Distortion**		
		dBc	N/A	
		% effected subcarriers	N/A	
	7	**Other Bands**	TBD	New Upstream spectrum
		dBc	-50	Note 3
		% effected subcarriers	1	50 kHz subcarriers
Wideband	8	**Burst Interference**		Note 4
		Bandwidth (MHz)	TBD	Non-white characteristics (Note 5)
		Level, dBc (PSD)	0	
		Duration (usec)	1	
		Period (Hz)	1000	
Freq Response				
	9	**Amplitude Slope**	N/A	Captured in Path Loss Range
Amplitude	10	**Amplitude Variation**		not included
		(dB pk-pk/24 MHz)	1.5	
		(dB pk-pk/96 MHz)	2.5	
		(dB pk-pk/192 MHz)	3	
Phase	11	**Group Delay Variation (nsec)**		
		Over 24 MHz		
		Mid Band	25	
		Band Edge (24 MHz)	280	
		Over 48 MHz		
		Mid Band	50	
		Band Edge (24 MHz)	305	
		Over 192 MHz	575	
Echo	12	**Echo Profile, dBc**		Note 6-7
		.5 usec	-16	
		1 usec	-22	
		1.5 usec	-29	
		2 usec	-35	
		3 usec	-42	
		4.5 usec	-51	
		5 usec		
Spurious Modulation	13	**AM/Carrier hum modulation**	5%	
	Notes			
	1	Path Loss adopted for consistency although return path, although RF actives include upstream gain		
	2	Measured samples in MSO location of high field strength environment		
	3	Projected (for 50 kHz) from acceptable D/S interference level for analog video band (now		
	4	U/S burst characterization in process; Ref CableLabs 1997 Report "Characterization of Upstream Transient		
	5	No linear optical return - no U/S Laser Clipping (white)		
	6	Measured Upstream CM (97% criteria) extrapolated to band (30 MHz measured to 100 MHz)		
	7	Echo mask range for a Single Dominant echo - does not imply an assumptions about multiple echoes		

CAN DOCSIS NETWORKS LEVERAGE SDN?

Gerry White
Arris

Abstract

Software Defined Networking (SDN) offers the promise of simplified network operation through centralization of functions, automation and programmable interfaces. Although SDN is still in the early phase of its evolution and has initially focused on data center networks it is appropriate to consider whether it could provide any advantages in a DOCSIS infrastructure. This is especially relevant as operators deliver more services from a cloud environment based on virtualized data center technology.

Accordingly this paper looks at the potential advantages of moving components of the DOCSIS ecosystem to an SDN environment. In particular it takes a detailed look at the CMTS and ancillary servers to determine whether a different decomposition of functionality would enable the industry to better leverage this technology. A possible implementation of an SDN based CMTS is proposed and compared to existing implementations As part of this evaluation the DOCSIS protocol itself is considered and potential changes suggested to remove roadblocks and make it more "SDN friendly".

INTRODUCTION

One of the most significant developments in service delivery over the last decade has been the evolution of data centers to deliver cloud computing on a massive scale. Low cost computing and storage platforms coupled with sophisticated virtualization techniques have enabled the delivery of complex services at very low cost points. The deployment of these large scale virtual server farms has illustrated some of the problems with conventional networking equipment, especially in the areas of cost and operational complexity. Software Defined Networking (SDN) has evolved as a potential answer to these problems. This paper looks at some of the basics of SDN, how it might be applied to deliver high speed data services in an MSO network, the benefits which could accrue and the problems to be resolved.

SDN

As workloads change in a data center large numbers of virtual machines must be instantiated, destroyed or moved to optimize use of compute and storage devices. The network must have the capacity to connect all of these devices at high speed with acceptable cost. More critically it must keep up with these changes as they occur which is problematic if new cables must be connected or even if large routing domains must re-converge on each topology change. Software Defined Networking has evolved as an attempt to address these challenges.

In a traditional network each network element such as a switch or router is composed of a data plane and a control plane. The control plane is typically composed of complex software components such as routing protocols executing in an embedded general purpose CPU on the device. The control plane software exchanges messages with other devices in the network (e.g. via a routing protocol) to determine the network topology and state and uses this information to create the forwarding tables used by the data plane to control packet forwarding. Thus each device has its own view of the network and all devices must cooperate to provide the required end to end paths. If forwarding is to work correctly then all of the local views must be coherent so that after a change or failure the devices must converge to a single view. This need for convergence causes a problem for rapidly changing topologies such as those found in data centers.

For high speed devices the data plane is typically built in hardware and performs line rate packet forwarding. When a packet arrives the data plane hardware looks up the address fields in the forwarding data base, selects the outbound port and forwards the packet.

This is of course a somewhat simplistic view of network operation but serves to illustrate the principals under consideration.

Figure 1 shows a sample SDN based network. It is composed of three layers, applications, network control and network infrastructure. The control plane functions which would be embedded in traditional network devices are moved to the control plane where they provide a centralized view of the network rather than the traditional distributed view. This single central view exposes Application Programming Interfaces (APIs) to the application layer. Thus business specific applications can control the network directly via high level interfaces rather than relying on the results of the traditional embedded control plane. The migration of the control plane away from the embedded devices also enables additional functionality to be provided. For example, the control plane can instantiate multiple virtual networks running over a single hardware infrastructure in a direct analogy to the virtual machine architecture used so successfully for data processing. The control plane itself can of course be run in a virtual machine environment to enable redundancy and scaling.

The network infrastructure contains the network devices themselves and provides the data forwarding plane for the network. The network devices in an SDN can be much simpler than traditional routers and switches as the control plane software has been removed and replaced by a much simpler module which updates the local forwarding database based on instructions from the central control plane. Thus the devices can be constructed from simple fast hardware with minimal software.

In order to enable the control plane to operate with network devices from different vendors a standard interface is required. Openflow [OPENF] is an interface specification defining the mechanism by which the control plane modifies the forwarding paths in the network devices. It is supported by multiple switch vendors and enables simple forwarding engines to be controlled by the SDN platform. Thus control of packet forwarding moves away from the network device and into the SDN platform. This enables low cost switches to replace more complex traditional switches and routers.

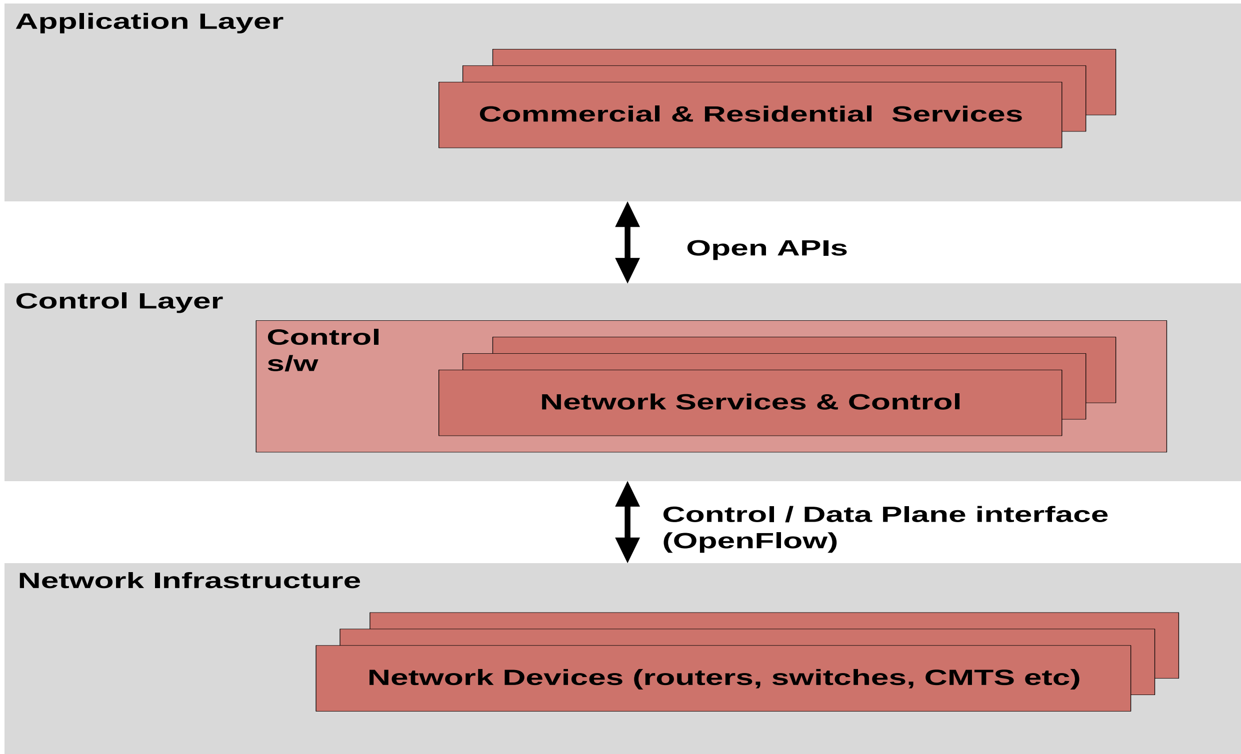

Figure 1 SDN Network Components

To date SDN has been primarily targeted at data center networks. The question to be considered is whether and how it could be used in a DOCSIS based MSO network and what benefits might accrue. The existing cable network is significantly different to those found in data centers. It is however evolving rapidly so that we will look at where it is likely to go and then look at the prospects for using SDN in this evolved network.

CHANGES IN CABLE INFRASTRUCTURE

Changes in cable infrastructure will be required to match the evolution of the services which must be delivered.

The following major trends are driving cable service evolution:

- Rapid expansion of the data rates which must be provided for high speed data services
- A move from broadcast to narrowcast video
- A move to IP delivery for video using Adaptive Bit Rate (ABR) protocols

The major impact of these changes to cable infrastructure is that they will require a significant expansion of the IP capacity currently in place. The existing core and regional networks are IP based but will need expansion to cope with the additional load. The access network is currently a hybrid of analog, MPEG and IP delivery and it is here that more dramatic changes will be required. As more services move to an IP base the capacity of the existing CMTS network used to deliver IP over DOCSIS will need to expand to accommodate this.

IP Expansion In the Access Network

Current CMTS ports are significantly more expensive than their video only MPEG equivalents. If this does not change then the dramatic expansion of IP capacity which will be needed could represent a significant CAPEX problem. This has been apparent to both operators and vendors for some time and has driven development of next generation DOCSIS platforms such as those based on the CCAP specifications [CCAP]. These platforms leverage high density silicon components to offer higher capacity and lower per channel costs in the same footprint as current CMTSs.

As the move to IP delivery continues it is appropriate to investigate whether additional changes to the MSO infrastructure would be advantageous and it is in this context that we will examine the potential for SDN.

Move to data center

Centralized data centers provide low cost processing for applications (and for the software components of an SDN). As described in [TRANS] a potential evolution of the cable network to leverage the advantages of a data center environment is a practical proposition. In this architecture the head end in its current form can be replaced by a data center, an Ethernet distribution hub and a simple node as shown in Figure 2. In this model all complex software has been moved to the data center and HFC specific MAC and PHY functions have been moved to the node allowing the use of standard Ethernet optics and switching from data center to node.

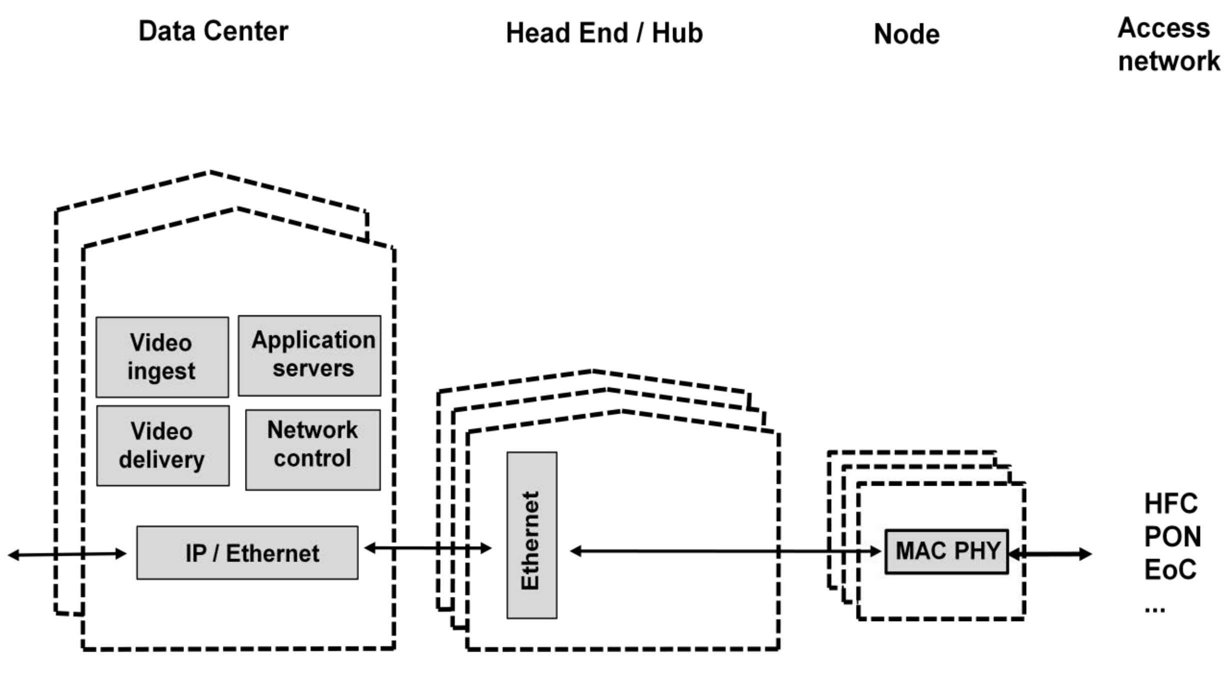

Figure 2: Distributed Cable Architecture

Moving the MAC and PHY functions from the head end to the node creates a more intelligent outside plant architecture. Operators who do not wish to take this step and prefer to keep a simpler outside plant can elect to deploy the MAC-PHY components in the hub rather than the node as shown in Figure 3. They still retain the advantages of the move to the data center and a significant reduction in hub complexity. Readers interested in an in depth comparison of traditional and intelligent HFC architectures are referred to [HFCDFC].

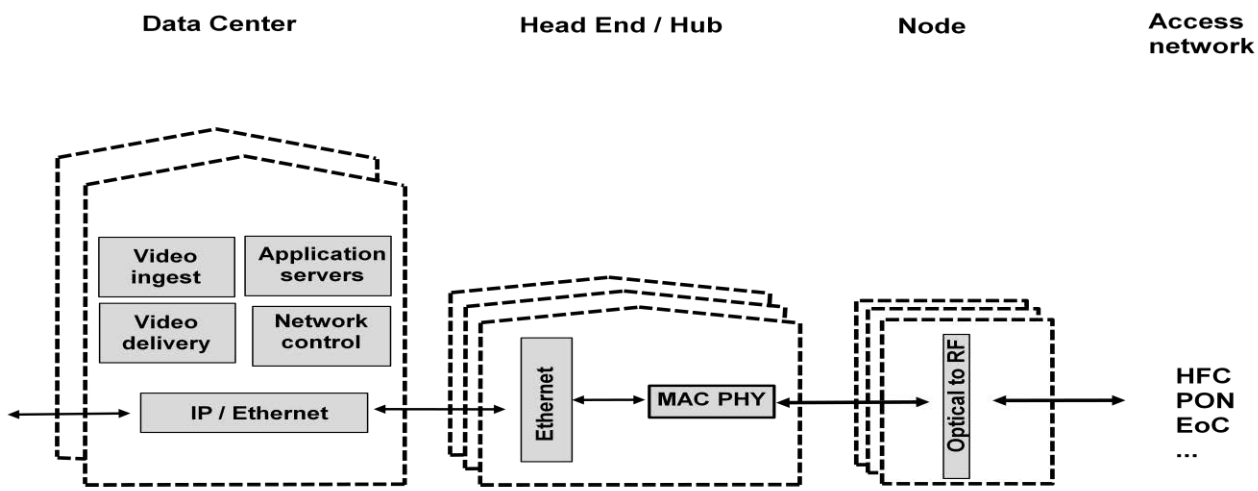

Figure 3: Passive HFC Architecture

SDN can be applied to either of these architectures in essentially the same manner. Thus they will be treated as equivalent for the purposes of this paper.

<u>DOCSIS HEAD END</u>

The previous sections outlined a potential move to a data center based architecture. We will now take a more in depth look at the DOCSIS head end infrastructure, how it could migrate to this type of platform and how SDN could be leveraged.

Figure 4 shows the major components of a PacketCable Multimedia system used to deliver QoS enabled multimedia services over DOCSIS. A detailed description can be found in [PCMM] but is not necessary to follow the paper which simply uses this as an example.

- An Application Manager /Server hosts the QoS-enabled applications and coordinates policy and QoS decisions.
- A Policy Server implements MSO-defined authorization and resource-management procedures which are enforced by the CMTS.
- The Record Keeping Server (RKS) tracks the usage of access-network QoS resources via message exchanges with the CMTS.
- A security server provides an authentication and security infrastructure support.
- The Operational Support system provides operations and management support.

- The CMTS provides data forwarding and DOCSIS control functions
- A managed IP network connects these devices together

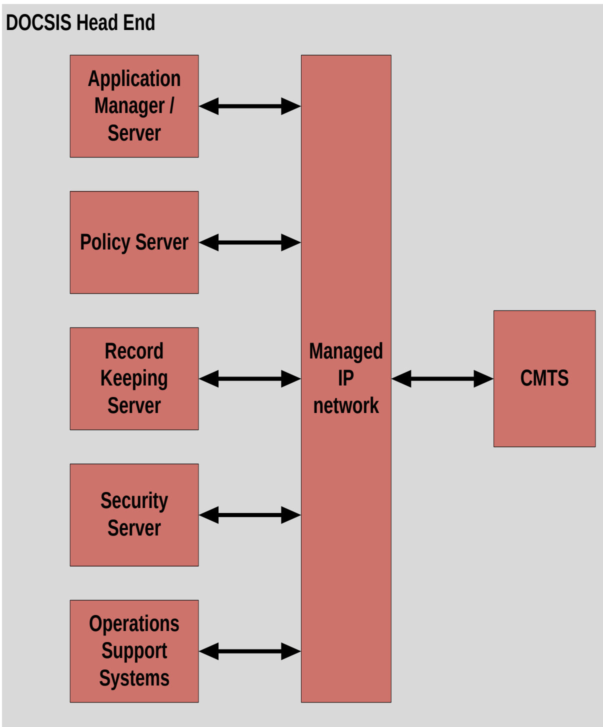

Figure 4 DOCSIS Head End

All of these components other than the CMTS and the IP network can be implemented as software running on standard server platforms and would be part of the application layer in an SDN based infrastructure as shown in Figure 5. The interfaces between the components would be modified for the centralized control layer but the functionality would remain largely intact.

The managed IP infrastructure could be evolved into an SDN architecture by moving the control plane into the SDN control layer and using OpenFlow switches in the same way that this has been done for data center networks.

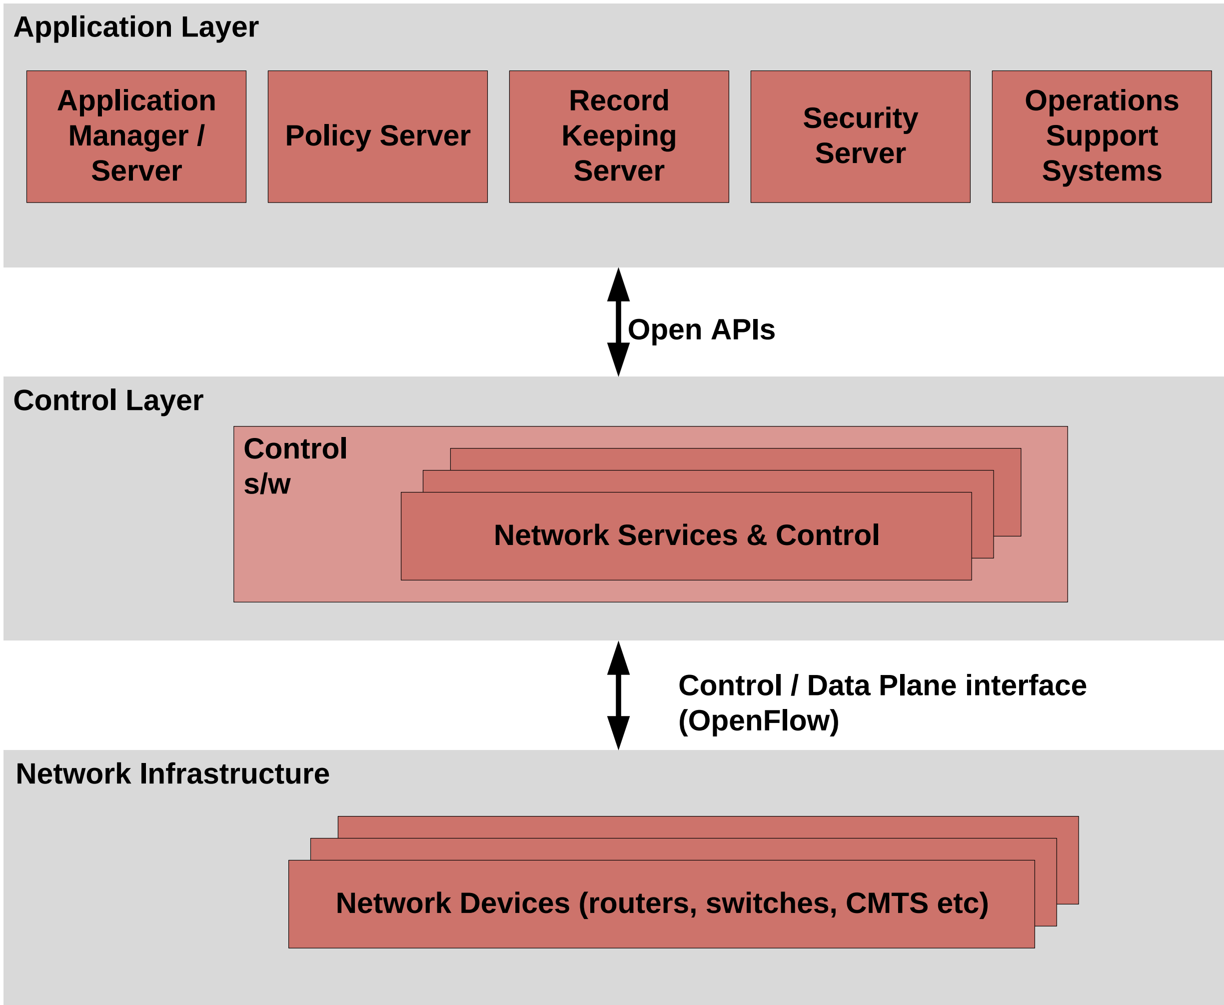

Figure 5 SDN DOCSIS Head End

CMTS

Moving the CMTS into an SDN architecture is a more complex undertaking which we will now examine in more detail.

Figure 6 shows the major components of an integrated CMTS system. These can be divided into three areas the control plane, the digital data plane and the RF data plane.

Control Plane

The control plane is typically implemented as software running on one or more general purpose CPUs and provides the following functions:

- Execution of the routing and layer 2 protocols used to create the forwarding data base
- Execution of the DOCSIS finite state machines used to control the interactions between the CMTS and the cable modems.
- Execution of the BPI+ security finite state machines used to manage the

security association between the CMTS and the cable modems.
- Interfaces to the control systems for PacketCable [PKCB] and PCMM used to establish QoS enabled sessions (shown in the application layer in Figure 5).
- Upstream bandwidth allocation and scheduling to manage the shared upstream resources
- DOCSIS specific control functions to handle CM operation and policy
- An operations and management component used for CMTS system provisioning and control.
- Additional application specific services such as subscriber management which may be present in some cases
- Slow path forwarding to handle exception packets which the data plane cannot forward e.g. ARP, DHCP.

Digital Data Plane

The digital data plane is the hardware based packet forwarding system and provides the following functions:

- Ethernet interfaces to connect to the regional and core networks
- The fast path forwarding engines which transfer packets between ingress and egress interfaces. Forwarding decisions are made by matching data fields in the packet header against entries in the forwarding database which has been created by the control plane. Packet header manipulation is also performed in this engine.
- The downstream QoS component implements the per flow QoS model defined in the DOCSIS specification.
- The DOCSIS MAC implements the lower layers of the DOCSIS MAC protocol in conjunction with the DOCSIS software components in the control plane. This includes DOCSIS header creation and removal and encryption / decryption functions.

RF Data Plane

The RF data plane provides the interface to the HFC plant:

- The downstream PHY provides conversion of digital input signals into QAM analog output, frequency conversion and RF transmission.
- The upstream PHY provides pre processing and demodulation of the received analog signals.

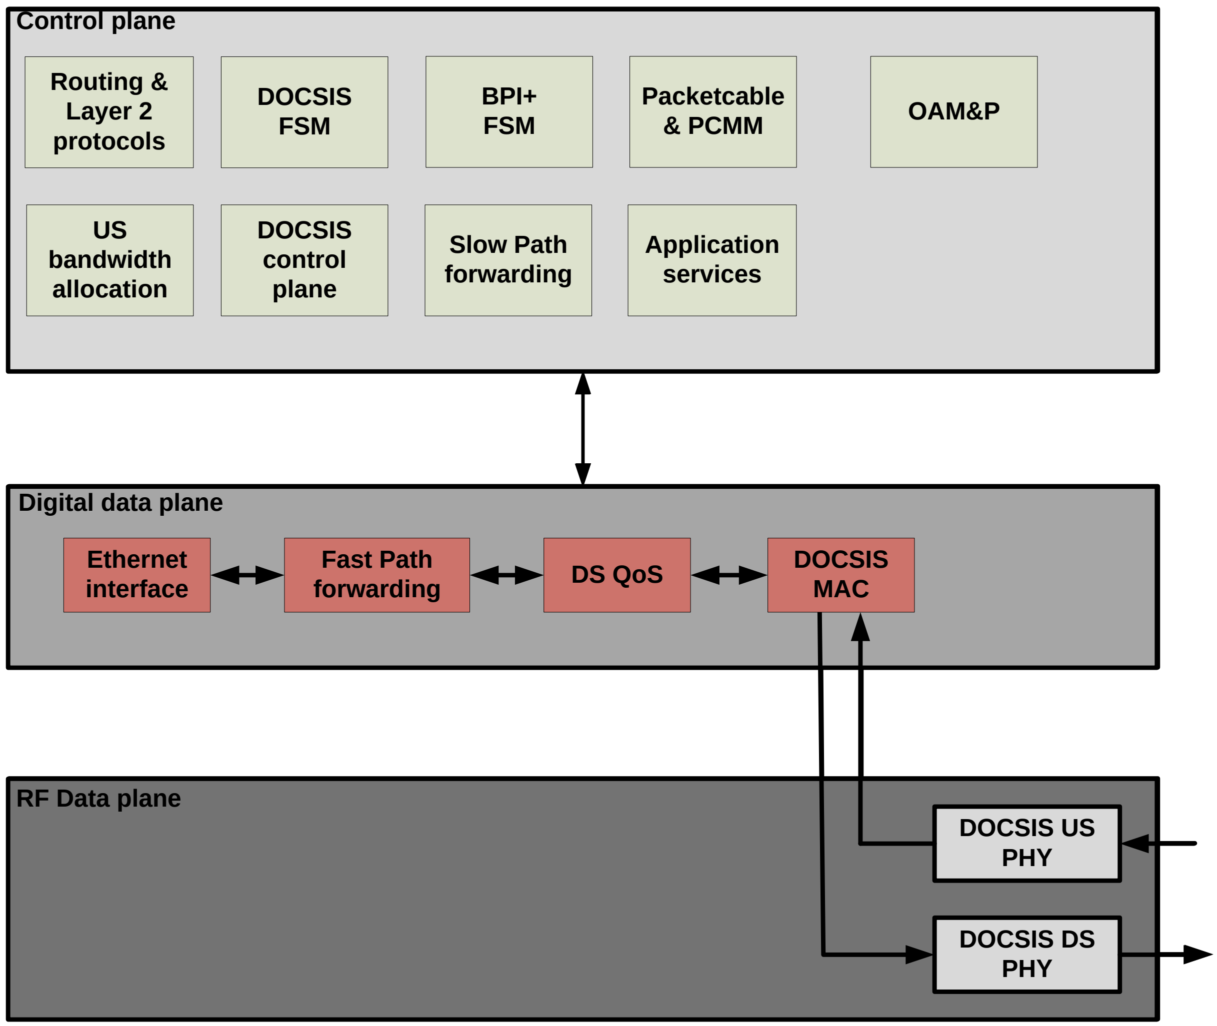

Figure 6 Integrated CMTS

A modular CMTS architecture [M-CMTS] as shown in Figure 7 is very similar except that the downstream PHY is implemented in a separate universal edge QAM and a DEPI [DEPI] control module is added to manage the interface between CMTS and UEQAM..

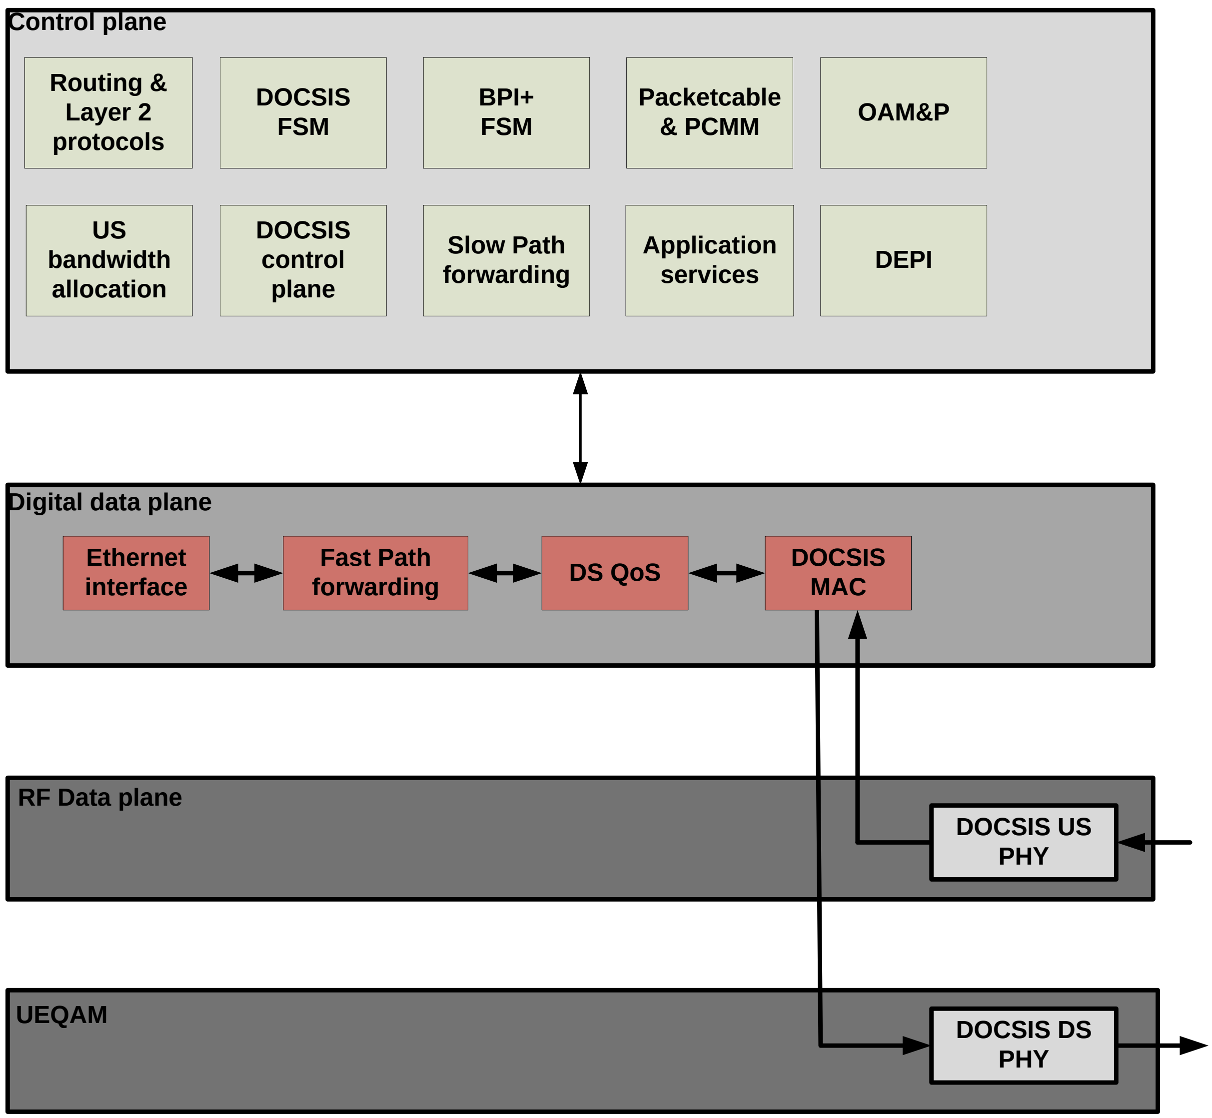

Figure 7 Modular CMTS

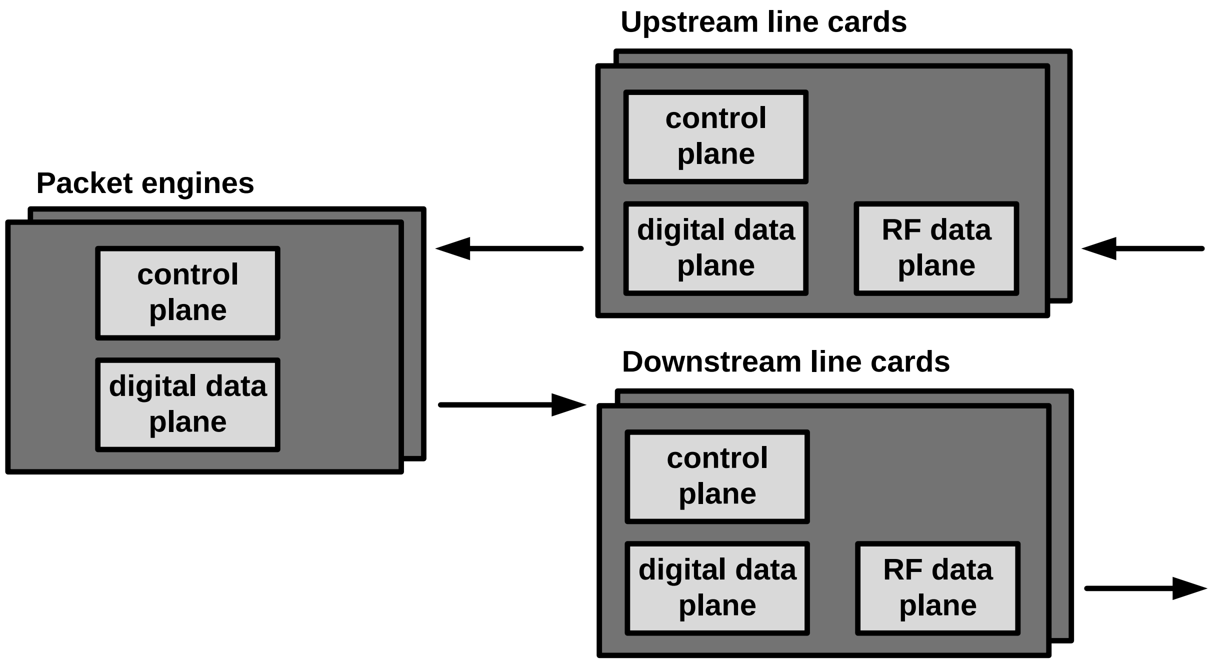

Figure 8 Typical Current CMTS Implementation

Figure 8 shows a typical CMTS implementation with the control plane distributed across all of the line cards in the system. In practice this distribution will be very uneven with the majority of the control plane software typically running in the packet engines.

SDN BASED CMTS

Figure 9 shows one option for how the CMTS can be moved to an SDN model. The control plane software is removed from the embedded control plane processors in the CMTS to the SDN control plane while the data plane of the CMTS remains essentially intact.

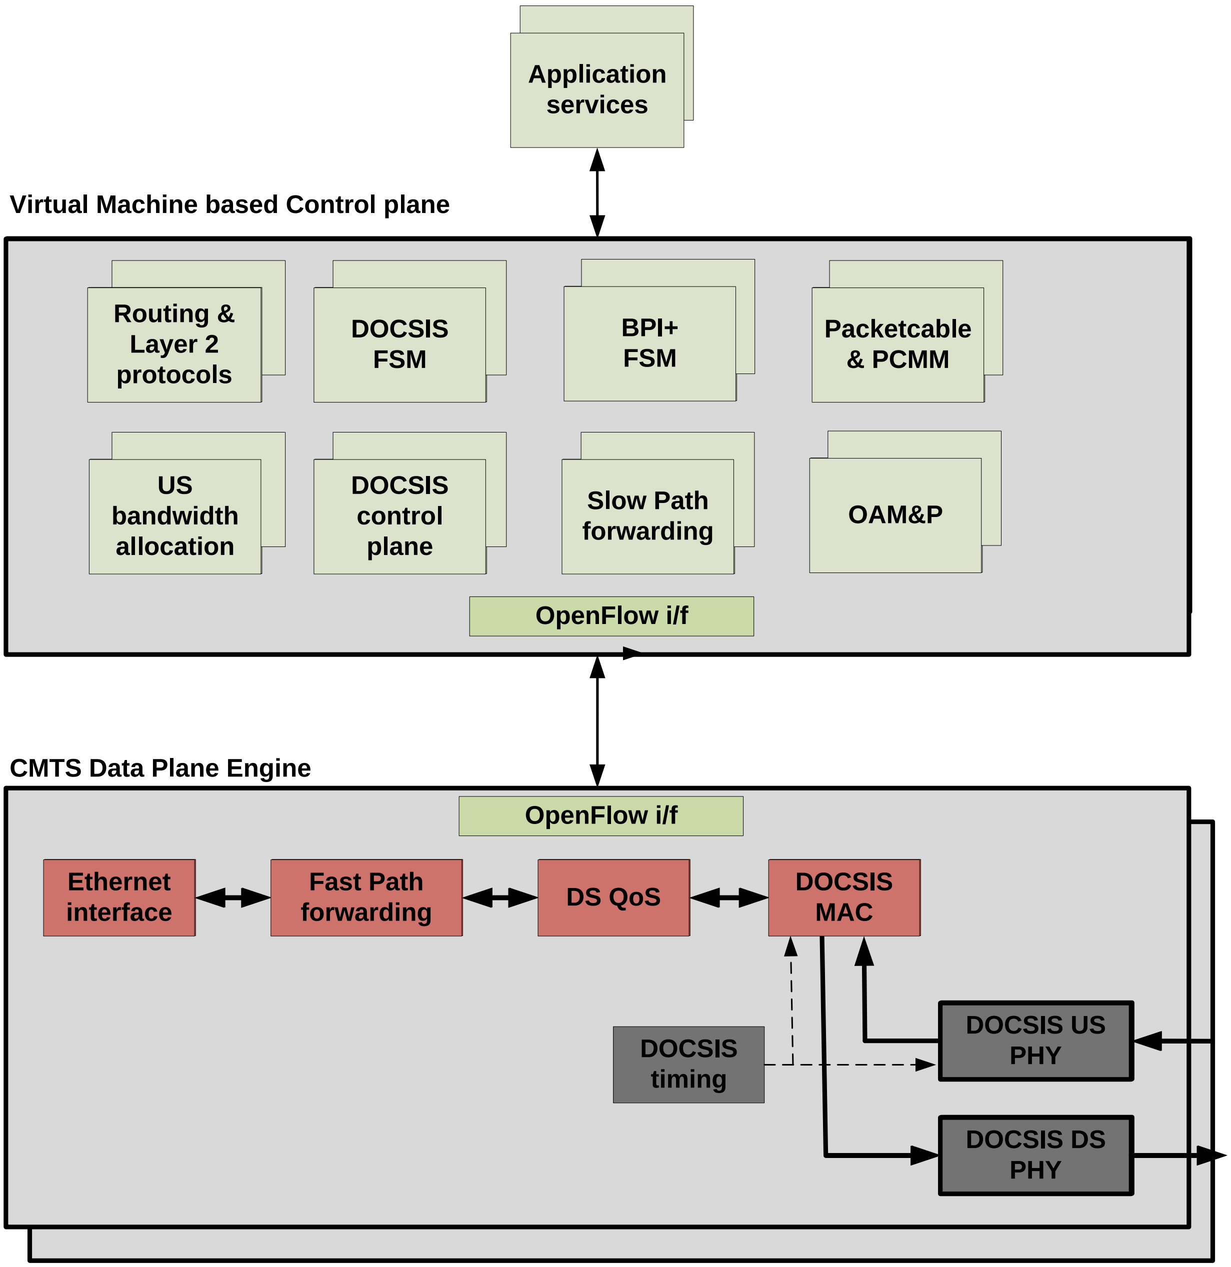

Figure 9 SDN based CMTS Phase 1

The control plane modules continue to provide the same functionality except that internal interfaces within the CMTS are replaced by messaging between the control and data planes. This would be based on the OpenFlow protocol.

- The routing and layer 2 protocols are likely be retained in an SDN system as it must interoperate with existing network devices such as the core routers. In this case they would use OpenFlow messages to update the forwarding data base in the data plane. With the move to a centralized control system it would of course be possible to construct a system in which the forwarding tables were constructed using a totally different mechanism without the need to change the OpenFlow based data plane devices.

- DOCSIS management messages continue to be exchanged between the finite state machines of the CMTS control plane and the cable modems via the data plane of the CMTS. They now incur an additional network hop between the CMTS data plane and the control plane so that some management of the added latency would be required.
- The BPI+ security finite state machines used to manage the security association between the CMTS and the cable modems will run in virtual machines in the control plane. In the SDN model the keying updates need to be passed to the encryption engines in both the CMTS and the CMs.
- PacketCable [PKCB] and PCMM systems running in the control plane VMS are used to establish QoS enabled sessions. They can continue to use the existing interfaces to the servers but could also move to more modern REST based interfaces.
- Upstream bandwidth allocation and scheduling are typically very CPU intensive. They can take advantage of the additional low cost CPU cycles available in the workstation hosted control plane. They can also be scaled up and down by adding or removing VMs as needed.
- The operations and management component used for system provisioning and control would run in the control plane. The configuration data for the data plane would need to be passed over the interface. Monitoring and alarm data need to flow in the opposite direction.
- Adding application specific services such as subscriber management becomes much more practical in the SDN environment. Additional CPU cycles are readily available on demand and the server based development environment is much friendlier than those for embedded systems.
- Packets which cannot be forwarded by the data path for any reason are handed off to software for processing. In a current CMTS this will run in one of the on board control plane processors. In the SDN case these packets will be handed off to the remote control plane for processing. The two inter system hops (data plane -> control plane -> data plane) added to the path will of course add latency but by definition ‘slow path’ packets, e.g. ARP or DHCP, need additional processing which adds latency so that this should not be a problem.

With this move we have taken a major step towards SDN, separation of control and data planes and centralization of the control plane functions.

Figure 10 shows The CMTS decomposition taken a step further. In this model the DOCSIS specific hardware has been separated from the IP/Ethernet packet forwarding engines and moved to a DOCSIS MAC-PHY shelf. Following this move the Ethernet interface, fast path forwarding and DS QoS modules are a very close approximation to an off the shelf Ethernet switch and we have taken a further step along the SDN path by enabling the use of COTS forwarding engines.

The DOCSIS MAC-PHY shelf is still needed to implement the DOCSIS framing and security functions at the MAC layer and for the DOCSIS PHY layer and RF functions.

DOCSIS QoS is based on a per flow model which is typically not supported in Ethernet switches. Thus if this must be supported an additional QoS module may be needed in the MAC-PHY shelf

The DOCSIS upstream and downstream links are tightly coupled so that a timing module is also required to ensure synchronization.

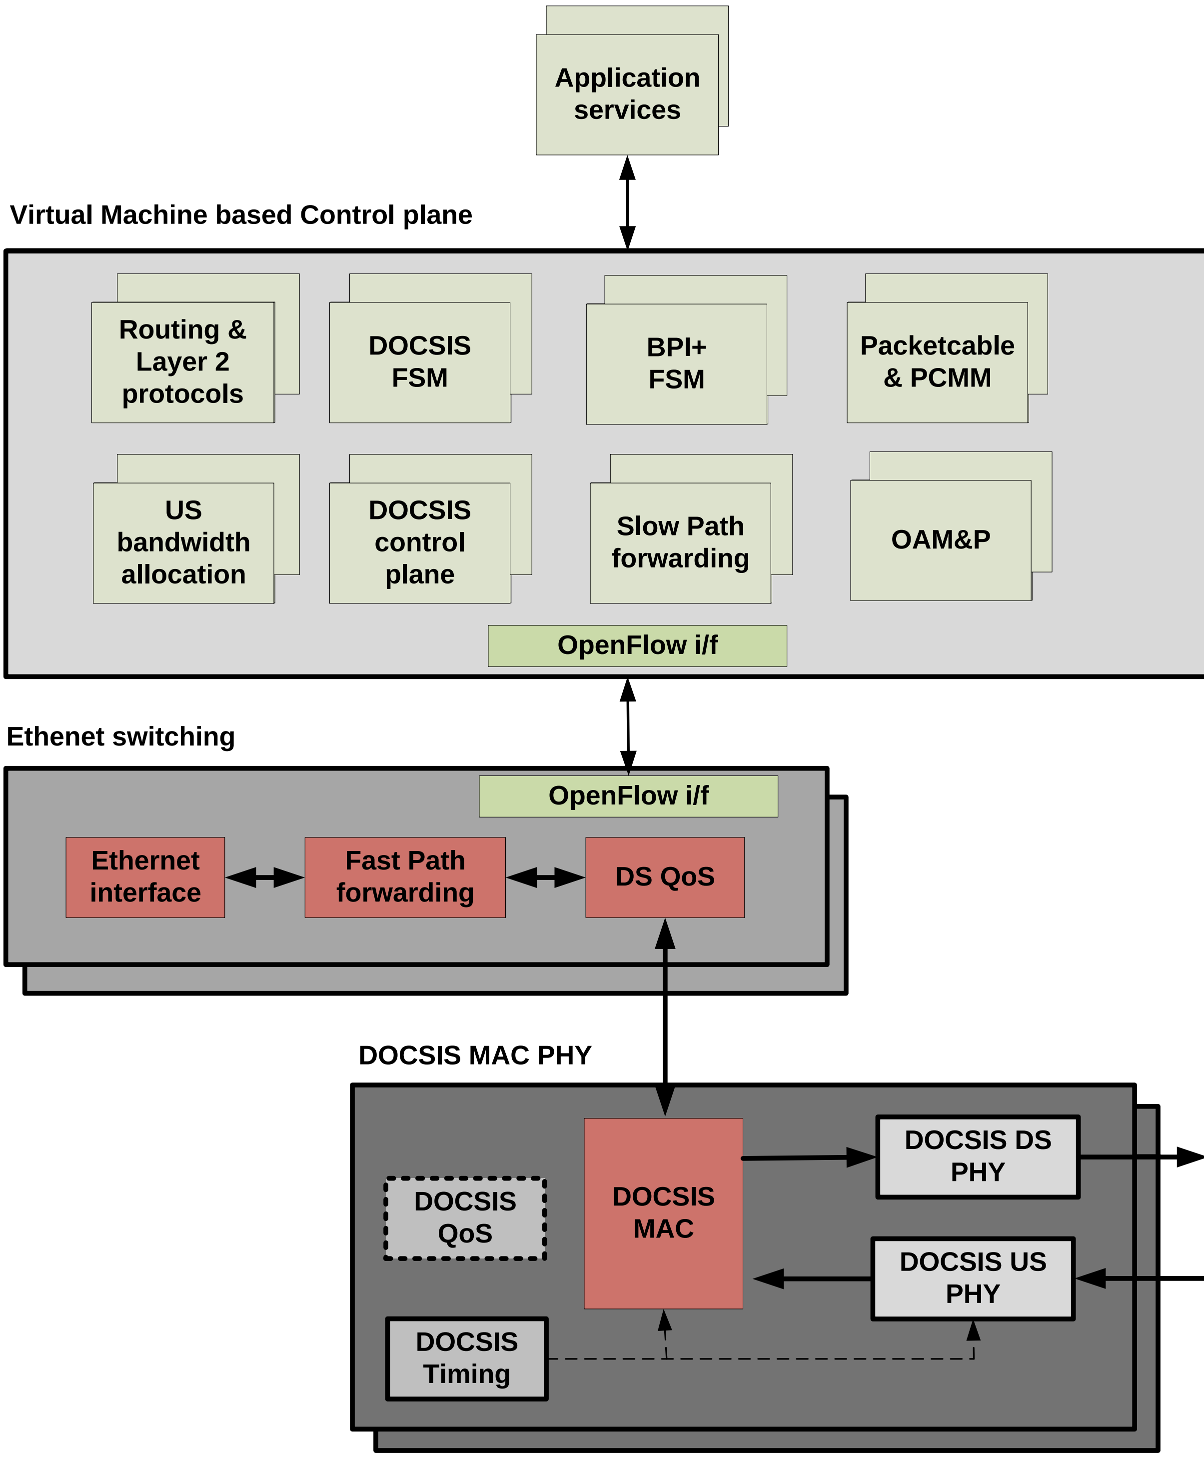

Figure 10 SDN Based CMTS Phase 2

SDN BENEFITS

As the cable infrastructure migrates to a data center architecture Software Defined Networking (SDN) can offer significant advantage in both network architecture and operation.

Centralization

The separation of control and forwarding in network devices such as routers, switches and CMTSs allows the control plane to be centralized rather than distributed across multiple devices. This has the benefit that it can provide high level APIs to application services to enable control of the network. This enables the network operation to be driven directly by the policies and needs of the applications rather than using policy to influence the operation of the distributed control plane as in the current model.

The control plane software in modern communications equipment is typically implemented as software running on top of a commercial RTOS such as Linux and executing in general purpose cpu's embedded into the system. This has the advantage of simplicity as the systems are self contained. The embedded nature of the devices does however have the following disadvantages:

- Embedded CPU cycles are very expensive compared to processing costs in general purpose servers. They also have a much slower upgrade cycle. To take advantage of newer (faster, cheaper) CPU versions these must be designed into the next generation of line cards and the deployed systems upgraded. Upgrading to new server hardware is a much simpler operation as the control plane VMs are migrated to the new server platform.
- Embedded software is more complicated to develop and test than workstation software due to inferior tool chains and more expensive test equipment. Thus development cycles are longer and costs higher. It is also more difficult to leverage third party software in an embedded system.
- The distributed control plane entities must interwork to create a coherent view of the network through complex routing protocols which need time to converge following any failures or topology changes

COTS Components

Moving to an SDN architecture enables the use of standard low cost compute platforms for the control plane and standard low cost forwarding engines to provide a significant portion of the data plane. This allows the CMTS hardware costs to benefit from the much larger scale of the enterprise market and allows multiple suppliers to participate.

Load Balancing

Horizontal scaling and redundancy through load balancing techniques is an established tool used in large scale data processing operations. The workload is distributed across multiple virtual machines which can be running on one or multiple physical servers. Figure 11 shows a simple example of multiple RTOS systems running as virtual machines with a VM hypervisor acting as arbiter to the physical hardware. Thus any failure of either software or hardware reduces the overall processing capacity but does not stop a function from executing.

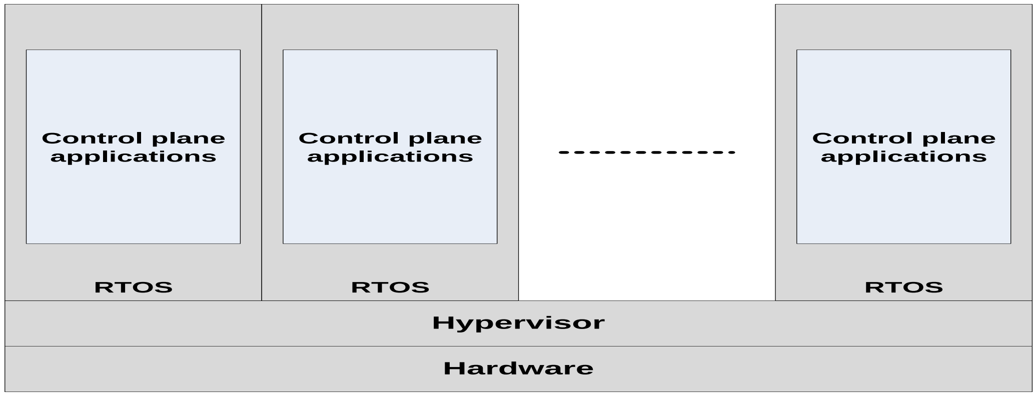

Figure 11 Control Plane Virtual Machines

Moving to this model has potential advantages over an existing CMTS. An RTOS partition can be be assigned to support a much smaller subset of the total user population than would a typical CMTS line card. A software failure would only impact this subset of users rather than all users on the line card. Essentially we use the hypervisor technology to scale the system and reduce the size of the software failure group. As the number of users in the system changes the number of VMs instantiated can be increased or decreased accordingly.

Upgrades to software are significantly simpler in the VM environment. A VM can be spun up with the new software version and tested against a subset of the users. As VMs are removed and restarted over time the upgrade can be managed and deployed from the central location. This is obviously a much simpler operation than upgrading large numbers of embedded systems distributed over multiple field locations.

Virtual Networks

SDN enables the creation of virtual networks which overlay the physical network infrastructure. The virtual network can be targeted at specific applications and manipulated via software changes rather than moving boxes and cables. It is a close analogy to the virtualization which has occurred in the server farms in data centers. The virtual network operating system is analogous to a virtual machine system such as VMWare. In a cable environment virtual networks could be created for residential high speed data, business services and managed IP video. Figure 12 shows a possible virtual network system based on OpenFlow interfaces to the Ethernet switches and DOCSIS MAC-PHY shelves

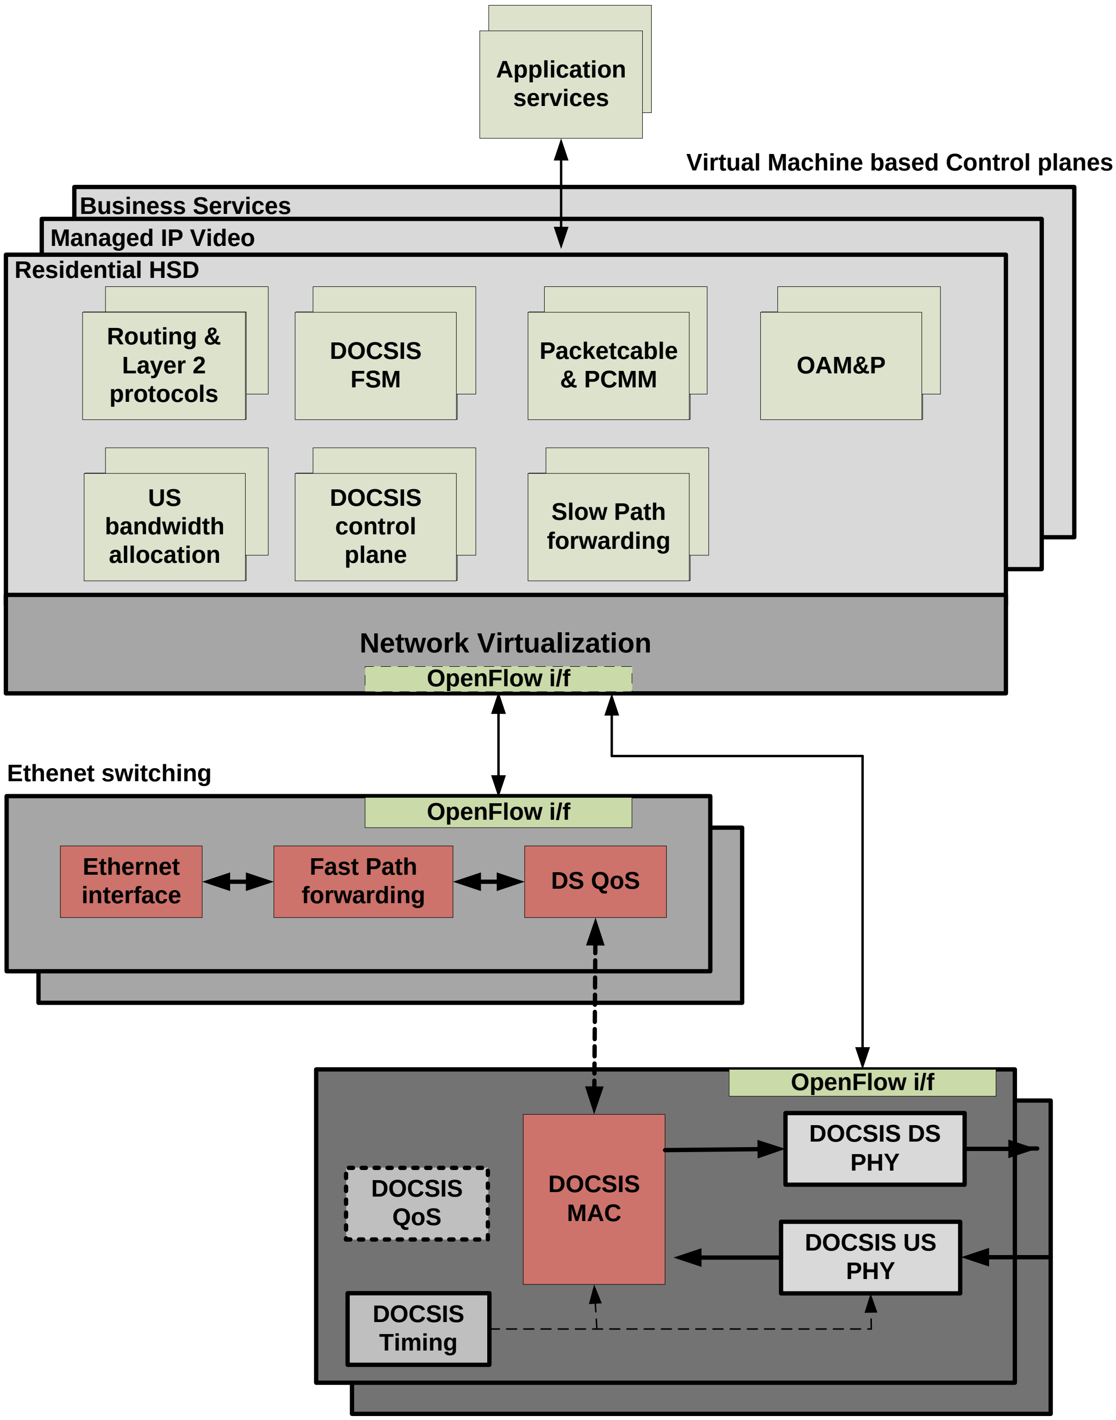

Figure 12 Virtual DOCSIS Networks

SDN PROBLEMS

Not all DOCSIS functions can be performed using off the shelf hardware components. The two suggested implementations of an SDN based CMTS Figure 9 and Figure 10 both require DOCSIS specific PHY, MAC and QoS hardware components. Thus a DOCSIS shelf will be required in addition to the servers and switches. The reasons for this are described below.

DOCSIS PHY

In the downstream direction the DOCSIS PHY hardware modulates the digital data stream into QAM signals, up-converts these to the desired frequency, conditions and delivers the RF signal to the network. The upstream PHY receives the RF signals from the cable modems and demodulates them to produce a digital byte stream. It is also responsible for power and timing measurements used in the ranging process [DOCSIS].

DOCSIS MAC

In the downstream direction the DOCSIS MAC hardware implements the lowest layers of the MAC protocol including DOCSIS header creation, packet encapsulation and content encryption. In the upstream direction decryption and header removal are performed.

DOCSIS QoS & Bonding

DOCSIS provides a very sophisticated per service flow QoS model enabling services such as voice and video to be delivered with guaranteed quality. The downstream QoS engine is more complex than the QoS available in a standard Ethernet switch so must be applied in the DOCSIS MAC-PHY shelf. With DOCSIS 3.0 channel bonding downstream QoS is even more complex as a packet may be sent on any one of multiple channels.

In the upstream QoS is provided by the upstream scheduling software running in the SDN control plane which allocates upstream bandwidth between requesting cable modems. For this to operate effectively it should not introduce significant added latency to the process, so that in practice this function may need to be split between entities in the DOCSIS shelf and the VM control plane.

DOCSIS Timing

DOCSIS requires that upstream and downstream links are synchronized to a common clock to enable efficient demodulation of the upstream traffic bursts. This timing function is restricted to the components in the DOCSIS MAC-PHY shelf and does not propagate to Ethernet switches or the control plane.

DOCSIS CHANGES

As shown above DOCSIS can be implemented in an SDN architecture but will still require some specialized hardware. Clearly some of this hardware such as the PHY components will always be needed but it is worth examining the protocol itself to see if it could be made more SDN friendly.

QoS

DOCSIS is primarily used to deliver Ethernet packets from servers to clients in a home network via HTTP connections. These packets travel over best effort (or possibly priority based) networks from the source to the CMTS and from the cable modem to the client with DOCSIS QoS being applied over the CMTS to CM link. Moving to a much simpler priority based QoS model would enable standard Ethernet switches to be used and remove the need for the downstream QoS hardware. This would potentially reduce the quality of PSTN style voice services but would still support VoIP and streaming video delivery, both of which currently operate successfully over cable networks as best effort services from OTT providers.

Bonding

DOCSIS 3.0 channel bonding allows multiple DOCSIS channels to be combined to offer a higher speed downstream service to a cable modem. Packets sent over the bonded link can arrive out of order so that sequence numbers must be added and checked. A bonding component is required on the MAC-PHY shelf to implement this. A DOCSIS bonding group can be created from any number of channels (up to the maximum supported) and multiple bonding groups can overlap almost arbitrarily. This makes packet scheduling a complex operation. If bonding groups were restricted to set numbers of channels e.g. 2,4,8,…) and constrained to be hierarchical this could simplify the bonding operation without a major loss of functionality and reduce the complexity of the MAC-PHY shelf.

Payload Header Suppression

PHS is a compression scheme introduced for the low bandwidth channels available in early DOCSIS deployments. With the move to much wider channels in the post DOCSIS 3.0 era this is rarely used and could be removed to simplify both embedded and SDN based implementations with minimal impact.

CONCLUSION

As the cable infrastructure migrates to a data center architecture Software Defined Networking (SDN) has the potential to have a major impact on network architecture and operation.

It can provide the following advantages over a traditional CMTS deployment:

- Separation of control and forwarding in network devices such as routers, switches and CMTS.
- The use of standard low cost compute platforms for the control plane
- The use of standard low cost forwarding engines to provide the data plane
- Centralization of control decisions as opposed to the decentralized model used by current network hardware
- Orchestration of resources across layers and domains for optimal performance
- Horizontal scaling and redundancy through load balancing techniques
- Provide general purpose interfaces to hide details of network elements from upper layer protocols and applications enabling simpler application development.
- The ability to provide virtual networks under the control of centralized applications.
- Remote operation with the control plane running in a cloud environment or super head end complex.

A number of DOCSIS protocol changes could make an SDN implementation simpler and lower cost. These should be considered for a future version of the protocol.

These advantages combine to offer the potential for lower capital and operational costs.

REFERENCES

[CCAP]	J. Salinger, "Proposed Next Generation Cable Access Network Architecture", SCTE Conference on Emerging Technology, 2009.
[TRANS]	Gerry White Transforming Cable Infrastructure Into a Cloud Environment, NCTA 2012 Spring Technical Forum
[HFCDFC]	M. Emmendorfer, S. Shape, T. Cloonan & Z. Maricevic Examining HFC and DFC (Digital Fiber Coax) Access Architectures, SCTE Cable -TEC Expo 2011
[M-CMTS]	Cablelabs DOCSIS® Specifications — Modular Headend Architecture (MHA)
[OPENF]	www.openflow.org
[DOCSIS]	Cablelabs DOCSIS® Specifications - www.cablelabs.com/cablemodem/specifications/index.html
[PCMM]	Cablelabs DOCSIS® Specifications - www.cablelabs.com/packetcable/specifications/multimedia.html
[PKCB]	Cablelabs DOCSIS® Specifications - www.cablelabs.com/packetcable/specifications/index.html l
[DEPI]	Cablelabs DOCSIS® Specifications - www.cablelabs.com/cablemodem/specifications/index.html

ABBREVIATIONS AND ACRONYMS

API	Application Programming Interface
CCAP	Converged Cable Access Platform
CDN	Content Delivery Network
CM	Cable Modem
CMTS	DOCSIS Cable Modem Termination System
COTS	Commercial Off The Shelf
CPE	Customer Premise Equipment
CPU	Central Processing Unit
DOCSIS	Data over Cable Service Interface Specification
FSM	Finite State Machine
Gbps	Gigabit per second
HFC	Hybrid Fiber Coaxial system
HSD	High Speed Data; broadband data service
HTTP	Hyper Text Transfer Protocol
IP	Internet Protocol
MAC	Media Access Control (layer)
Mbps	Megabit per second
MSO	Multiple System Operator
OTT	Over The Top
PHY	Physical (layer)
PMD	Physical Medium Dependent (layer)
REST	Representational State Transfer
RF	Radio Frequency
RTOS	Real Time Operating System
SDN	Software Defined Networking
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VM	Virtual Machine
VoIP	Voice over IP

DISTRIBUTED DIGITAL HFC ARCHITECTURE EXPANDS BI-DIRECTIONAL CAPACITY

Oleh Sniezko, Chief Technology Officer
Doug Combs, Vice President, Network Architecture & Field Services
Rei Brockett, Product Manager, Video Products
Aurora Networks, Inc.

Abstract

This past year has brought about a re-kindling of interest in the concept of moving the boundary between digital content (both data and video) and the RF domain away from the headend and further into the HFC network, out to the fiber optic node. While few dispute certain operational benefits, such as the ability to change QAM allocations more easilty, or the ease of "set and forget" installations, there is still a great deal of debate around what functions should be moved, in what order, and in what combinations. This paper attempts to clarify some of the tradeoff decisions by presenting a careful analysis of the performance profiles of today's analog forward and reverse optical links, and contrasts them with the gains enabled by transitioning to digital links. This paper also quantifies the benefits, using digital re-designs of actual N+5 and N+0 HFC systems.

INTRODUCTION

Since the early 1990's, the cable industry has invested heavily in building the hybrid fiber coax (HFC) infrastructure that is unequaled today in its ability to deliver bandwidth-intensive services to hundreds of millions of subscribers worldwide. A perennial question that haunts the industry is whether the HFC architecture is a viable solution for the next twenty years of evolution in subscriber demands. The simple answer is an emphatic "Yes!"

The traditional HFC network (Figure 1) comprises five major elements contributing to the system impairments. Some networks have fewer elements. For example, Fiber Deep N+0 networks do not have RF amplifiers after the node, which is critical to enabling forward bandwidth expansion, and also reduces return path noise funneling. Similarly, RFoG networks move the fiber-to-coaxial cable interface to the home, and significantly lower CPE contribution to signal performance degradation by providing relatively high signal levels for in-house wiring and networking.

There are complementary layers of innovation at the cable headend to support increasing levels of spectral efficiency for any given signal level, and at the customer premises to control in-home signal degradation. The key to promoting the HFC architecture's resiliency, however, is in reclaiming the signal performance margin that is currently consumed by the outside plant in large part due to the complex art of propagating RF signals over fiber (a.k.a. "analog fiber").

Technology and subscriber demands have combined to create an industry-wide push to initiate a new phase of growth, significantly expanding forward and reverse capacity above current levels. It is therefore prudent to perform a thorough review of optical link options between headends and nodes.

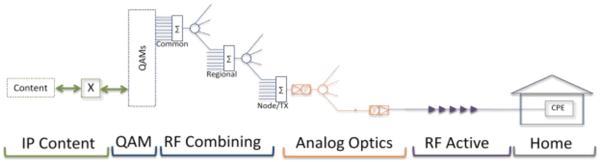

Figure 1: Five Major Elements of Traditional HFC Networks

EXISTING NETWORK ANALYSIS

Analog Forward Optical Links

Analog optical links (analog fiber) have served the cable telecommunications industry since late 1980's and have been crucial to the success of the broadband HFC telecommunications networks. They provided significant improvements in network reliability and service availability while maintaining optimal performance for signals carried on coax. They also significantly lowered operational costs and enabled network simplification by:

- Eliminating AML and FML links, which were used to serve remote pockets of subscribers but were burdened with regulatory and licensing compliance and subject to weather-related signal fading;
- Reducing RF amplifier cascades, which contribute significantly to noise and reliability issues, especially in the upstream path;
- Consolidating headends and hubs.

Analog optical links supported centralized signal processing at the time when the technology for distributed signal processing in the nodes was not practical without installation of secondary hubs and sizeable OTNs. With advances in silicon, especially in FPGA technology, distributed signal processing in the nodes is increasingly cost-competitive with and as reliable as traditional centralized processing. It is now possible to blur the functional boundaries between headend and node. As technology progresses, the choice of where to place the demarcation point is increasingly a question of operational organization rather than the feasibility of physical or functional capabilities in the node.

With the progress of the technology and the need for sizeable capacity gains in upstream and downstream, driven by demand and competition, it is important to understand the sources and causes of signal degradation in existing analog links, especially in relation to the content payload that is transmitted within the links.

OMI: Composite and Spectral Density

In a traditional, centralized architecture, signals are carried in their final RF format from headend through analog fibers and coaxial distribution network to CPEs. Hence, analog fiber OMI limits apply, and the performance of analog fiber is cascaded with the performance of all other network components from the RF signal source to the CPE.

In addition, the usable frequency bandwidth of fiber links in multi-octave HFC networks is constrained by the RF amplifier technology and cascades when carrying an RF signal load.

The original analog fiber transmitters were designed with sufficient OMI to support analog NTSC video in combination with 64- and 256-QAM signals over 450 MHz plant. As plant bandwidth expanded to 550 MHz, then 625 MHz, 750 MHz, and now 1 GHz, the composite OMI was distributed over increasingly wider bandwidth, leading to an increasingly lower OMI/Hz and OMI/channel. This trend is exacerbated by the discontinuation of analog NTSC video. In the presence of large numbers of analog NTSC channels, with their significantly higher CNR requirements, bandwidth expansion for QAM channel load added at levels 6 dB lower resulted in a very limited drop in OMI/Hz.

But in HFC networks with dominant QAM loading and even distribution of RF levels, the drop in OMI/Hz is noticeable. This results in per-channel degradation of analog fiber performance.

In HFC networks with RF amplifiers, the bandwidth capacity can be expanded by improving the end-to-end SNR, up to the capacity limit defined by Shannon's law:

$$C=B \log_2 (1+S/N)$$

With network CNR (SNR) of 33 dB (for 256-QAM signals), the theoretical capacity would reach approximately 11 bits/Hz and close to 9 Gbps in 154 6 MHz QAM channels. The real capacity with 154 256-QAM channels approximates 6.2 Gbps, remarkably close to the theoretical limits in the real network with all its operational margin requirements and changing operational conditions as well as equipment implementation losses. This is an endorsement for HFC networks, their design rules, and the operational practices developed by the cable industry and implemented by the network operators.

This capacity can be further increased approximately 15% to 20% (close to 7.5 Gbps after accounting for overhead losses) by increasing the FEC power . Further capacity expansion would require more efficient modulation/coding schemes. OFDM could increase the capacity by an additional 10% to 15% by limiting guardbands between RF digital carriers (increasing effective BW).

In the HFC network designed to carry 79 analog NTSC carriers (or equivalent analog load – minimum 50% of the operational bandwidth), some spectrum segments support better EOL performance (at least 10 dB better). These segments can carry digital RF carriers with higher modulation levels. Table 1 presents possible forward frequency allocation schemes for a 1 GHz HFC network and its related capacity. It also presents the increase in potential capacity for a traditional (not distributed) HFC N+5 network subject to a limited upgrade.

	Bandwidth (Channel Allocation)				
Bandwidth	Analog Channels	Priority Load Adv. PHY Digital RF Channels (BW)	Advanced PHY Digital RF Channels (BW)	Legacy Digital RF Channels	Approximate Total Data Capacity
54 to 1002 MHz	64	0	0	90	3.6 Gbps
	64	0	192 MHz	58	4.1 Gbps
	32	192 MHz	192 MHz	58	6.1 Gbps
	32	192 MHz	384 MHz	24	6.3 Gbps
	0	0	0	154	6.2 Gbps
	0	384 MHz	384 MHz	24	8.4 Gbps
282 to 1050 MHz	0	384 MHz	384 MHz	0	7.4 Gbps
282 to 1194 MHz	0	192 MHz	576 MHz	24	8.0 Gbps
	0	192 MHz	720 MHz	0	8.4 Gbps

Table 1: Enhanced Network Capacity (up to 1 GHz and above). Example for North America

As apparent from the table, in 1 GHz HFC networks designed to carry 79 analog NTSC channels and 75 RF digital channels, significant capacity gains can be realized by replacing analog channels and some existing RF digital channels (subject to the need to support legacy services) with RF digital channels with more robust FEC and OFDM coding.

This option exists for any network which has been designed to carry analog cable channels. However, in networks designed for a reduced upper operational frequency, the resultant increase in capacity would be lower. Moreover, in networks designed for QAM-only load, the capacity gains would be limited because there are no spectrum segments designed for higher performance; their design has been cost-optimized to reliably support services for 256-QAM signals as defined in ITU J-83, but some gains can be realized by replacing the existing QAM channels with RF digital channels with more robust FEC and OFDM coding which can be carried on the network with performance which supports the carriage of 256-QAM signals.

Further capacity gains would require either sizeable improvement of the EOL performance or bandwidth expansion, or both.

Operational Margins: The Reprieve Granted by VSB Modulation

The traditional North American HFC network loading contains 79 analog NTSC modulated carriers, which closely resembles the ratio of analog to QAM channels in other regions. In 1 GHz operational HFC networks, these analog channels are usually accompanied by 75 QAM modulated digital carriers. The 256-QAM signals are set usually 6 dB lower in RMS power relative to the equivalent RMS peak value of NTSC channels (or RMS values of CW carriers used during the testing). The analog carriers are considered "priority loading" due to their higher CNR requirements (relative to the CNR requirements for 256-QAM carriers as defined in ITU J-83).

The optical link OMI levels are optimized to allow for acceptable contribution of CTB and CSO and optical link noise (from all noise contributors) based on performance allocation among different HFC network elements (different for different operators and architectures) while still maintaining acceptable margin to QAM channel clipping. The commonly accepted clipping margin is set for 1 dB from BER values of 10E(-5) while a 1 GHz capable analog optical transmitter is tested with the nominal load of 79 CW carriers and 75 QAM carriers (either Annex B or C for 6 MHz wide channels or Annex A for 8 MHz wide channels).

In the operating environment, CW carriers are modulated with an average power 4 dB lower than their peak power (and with 79 channels, this lower power is realized). This margin is critical in an operational environment which is subject to test equipment errors, level setup errors and signal fluctuation, and ultimately results in quite reliable operation of analog optical links without clipping. It also changes the operating point of the laser transmitter. To maintain this reliable operation, it is recommended that this operational margin is maintained with the load that comprises exclusively RF modulated digital carriers or that the operational inaccuracies and variabilities are remedied at the analog fiber transmitter with composite power ALC to avoid clipping. Even in that case, 1 dB additional operating margin is recommended for reliable operation of the network with only an RF digital carrier load.

Table 2 compares the relative composite OMI change due to this modulation versus OMI with CW and QAM channel load. This table also presents the OMI/Hz change for different loads for two recommended operational margins.

Before we summarize the results of the analysis, the following assumptions and disclaimers should be noted :

1. The analysis presents only the results of relative OMI spectral density for fixed composite OMI values of the lasers and disregards RIN changes with frequency expansion.
2. It also assumes the same shot noise, receiver noise and IIN in the link as well as the same other impairment contribution in single- and multi-wavelength analog fiber links.
3. Consequently, it assumes that the analog fiber links were designed to provide adequate performance for the network carrying 79 NTSC analog channels and 75 256-QAM channels (or equivalent loads). For networks optimized to carry loads optimized for materially fewer analog NTSC channels or QAM-only load, the analog fiber links would have to be redesigned with:
   a. Transmitters with higher launch power (and the same OMI capabilities) assuming the launch power is not limited by other considerations;
   b. Transmitters with better RIN (especially if the receiver optical input levels are relatively high);
   c. Transmitters with higher OMI capabilities (it may affect other noise contributing mechanisms in analog fiber links);
   d. Receivers with much better performance, or
   e. A combination of the above remedies if they can be implemented.
4. If the analog fiber link transmitter was replaced, contribution from all other impairment sources and mechanisms must be reassessed.

Also, this analog fiber link performance improvement effort may be trumped by the BW expansion limitation of the RF coaxial section of the HFC system.

Traditional RF load		Digital RF carrier only load	
79 CW and 75 256-QAM signals	79 NTSC carriers and 75 256-QAM signals	Recommended with TX ALC	Recommended without TX ALC
[dB]	[dB]	[dB]	[dB]
-1	-3.9	-1	-3.9

Relative composite OMI referenced to maximal composite OMI for BER 10E(-5) of 256-QAM signals

	Existing Bandwidth 54 to 1002 MHz		
	154 256-QAM load	122 256-QAM & 192 MHz of priority load	90 256-QAM & 384 MHz of priority load
	[dB]	[dB]	[dB]
	924 MHz	732 & 192 MHz	540 & 384 MHz
TX with ALC	+3	+1/+7	0/+6
TX without ALC	+1	0/+4	0/+2.3

	Extended Bandwidth Analysis				
	186 256-QAM load	154 256-QAM & 192 MHz of priority load	186 256-QAM & 192 MHz of priority load	286 256-QAM load	314 256-QAM load
	[dB]	[dB]	[dB]	[dB]	[dB]
	1116 MHz	924 & 192 MHz	1308 MHz	1716 MHz	1884 MHz
TX w/ ALC	+2.6	+0.4/+6.4	0/+6	+0.4	-0.5
TX w/o ALC	+0.3	0/+1.7	NA	-1.5	-2

OMI per channel or Hz relative to the OMI per channel or Hz for 79 CW and 75 256-QAM load

Table 2: Relative OMI Analysis of Analog Fiber

The results of the analysis show that the replacement of analog channels with RF modulated digital channels within 1 GHz operating BW (54 to 1,002 MHz) in analog fiber links with ALC TXs enables approximately 3 dB increase in OMI per RF digital channel for all channels at the same level. Two other examples of loads result in a 6 dB higher OMI/Hz for a single priority load of 192 MHz BW segment and two priority loads of 192 MHz BW segment while maintaining the same or better OMI/channel for 256-QAM channels as exhibited in the link with 79 analog and 75 RF digital channel load. This approach preserves the performance of QAM-signals. In summary, the analysis supports the numbers presented in Table 1. It assumes that the source of the digital signal in the headend has sufficiently higher performance and the RF combining network is eliminated for the new signals (in fact, even if it remained, its impact is minimal).

The analysis also shows that the expansion of the optical fiber link bandwidth to 1.2 GHz and 1.4 GHz is also supported (assuming no degradation from RF amplifiers) with the addition of a single 192 MHz priority OMI segment. Further, the analysis shows that it is possible for analog fiber links to support in excess of 1.8 GHz of bandwidth with performance equivalent to the performance of 256-QAM signals.

Whether this expansion can be supported by the coaxial part of the HFC network will be analyzed further.

It is an operator’s choice to maintain, increase or lower the operating margin to clipping recommended by the authors. However, with OFDM signals and their higher PAPR, the clipping is more likely to occur. Even if the OFDM signals are more immune to clipping due to their characteristics as well as symbol rate and duration, in a hybrid load system all signals are clipped at the same time so 256-QAM signals supporting legacy services will be clipped much more often in a load with OFDM channels than in a purely QAM load. The operational margin is even more critical unless PAPR reduction methods are applied to the OFDM signals.

One last note: improvements in performance in high-performance optical links have a lower impact on EOL performance because the HFC network design is cost-optimized to take advantage of higher link performance. HFC networks deliver performance at the customer outlet: 47 to 49 dB CNR in frequency ranges carrying analog TV channels and 37 to 39 dB CNR and MER in frequency ranges carrying 256-QAM signals. This fact limits bandwidth expansion options that could be realized by incremental improvements in performance of the analog fiber link (if possible) without significant upgrades in the coaxial part of the HFC network. It is in this case where a step performance improvement can be realized with the replacement of the analog fiber link facilitating bandwidth expansion with just limited upgrades in the coaxial section of the HFC plant.

Analog Reverse Links

The reverse link analysis will start with an explanation of the operating requirements of the network.

Dynamic range

The dynamic range of the reverse optical link is defined as the range of the input power from the level where the link provides sufficient performance (CNR) to support end-to-end transport of the signals to the level where the reverse transmitter introduces clipping resulting in 10E(-6) BER. The CNR requirements for the reverse optical link will depend on the network configuration (how many CPEs and RF amplifiers are funneled into a single transmitter), the required combining levels and the CNR contribution of the RF splitting network and RF signal receivers (e.g., CMTSs) in the headend.

Dynamic range depends on several factors:

- Long loop AGC level range (hysteresis) and the operating point of the CMs within their output level range,
- Optical level stability of the transmitter,
- Gain/loss stability of the headend RF splitting/amplification network,
- Funneling of ingress and other impairments.

The long loop AGC hysteresis depends on implementation and could be as wide as ±3 dB (the wider window is beneficial for some considerations but is detrimental to reverse transmitter load fluctuations). In this analysis, it is assumed that the CM operates within its range of output levels so it is capable of reducing power when directed by long loop AGC. The analog optical transmitter power changes can be quite dramatic but well designed transmitters stay within ±1 dB (for the best designs sometimes lower). This is equivalent to a ±2 dB RF level change on the optical transmitter input enforced by long loop AGC. The RF splitting network, barring some unintended human error, should be reasonably stable but it is prudent to assume ±0.5 dB change. These three elements add to a 5.5 dB dynamic range requirement (level swing above the nominal link setting point). The additional contributors (ingress and noise funneling) depend on the network configuration and can be as high as 4-6 dB for N+5 HFC network, 2-4 dB for FD network and very little for RFoG networks (a subject for a separate discussion).

This results in a requirement of approximately 10 dB dynamic range for the analog reverse fiber link. Under simplified assumptions such as similar quality receiver and others, this translates to 10 dB better link performance requirements for the same distance, signal performance requirements and the same BW as in the forward if the RF amplifier funneling noise is disregarded and reverse receivers are not combined or the same and lower link performance requirements from lowering load BW, signal performance requirements and/or reach (The latter should match the forward reach unless E-O-E regeneration is implemented.)

Broadband Digital Links

Digital links have several advantages over the analog fiber links. The most important are:

- Link performance does not change with distance within the specified range.
- The link is extremely thermally stable and hence its dynamic range requirements can be lowered by 2 dB. (The input level operating point can be raised by 2 dB thus improving link CNR by 2 dB).
- The link performance does not change with the design BW as long as the same number of coding bits is used and the ratio of the sampling frequency to the upper operational frequency is the same.

Of course, as with an analog link, the link performance depends on the ratio of the real BW load to the maximum designed BW load because for lower real BW load the spectral density of the input signal can be increased within the limits of the maximum composite input power for the defined dynamic range.

Hence, for digital reverse links (and forward too), the considerations are related to and the decision is based on the cost of the link and the status of the technology. To their advantage, digital links provide all the benefits of analog links, with one notable exception, and bring additional benefits while avoiding all their shortcomings. If cost is not an issue, digital links can be made arbitrarily transparent to any signal requirements.

RF Headend and Hub Combining Network

The headend combining network is a negligible contributor to the system performance today but with a higher performance requirement for higher modulation signals, it cannot be neglected going forward. The main degradation is caused by crosstalk between narrowcast signals through the broadcast combining network which was built for analog fiber links with full spectrum transmitters (this degradation is eliminated in BC/NC overlay systems).

RF Amplifiers

RF amplifiers are a necessary evil in HFC networks when an optical node serves a large area through long coaxial cable runs. They compensate for coaxial cable loss which increases drastically with frequency, and for RF passive loss (reasonably flat in the most recent designs for up to 2 GHz and beyond). RF amplifier technology has progressed over the years. Most recently, after the first deployment of HFC networks, the operational BW of RF amplifiers increased from 450 MHz to 1002 MHz and their composite power output capability by 8 to 10 dB with the introduction of the first power doubling hybrids. This means that for a cable span of 18 dB loss, progress in RF amplifier silicon compensated for the increase in cable loss from 450 MHz to 1002 MHz. However, for longer cable spans, the RF amplifier network would have to be re-designed (re-spaced) or other network elements contributing to the signal impairments materially improved or eliminated. Alternatively, the signals placed above 450 MHz would be more immune to cable impairments and hence require lower levels thus increasing the “equivalent” level of the amplifier output power. Typically, QAM signals above 550 MHz and now above 450 MHz are at a 6 dB lower level relative to the equivalent analog channel level. This allowed for significant slope increase to compensate for the cable loss not compensated for by a simple linear extension of the existing slope.

The most recent introduction of GaN hybrids with increased output power capability enables linear extension of the slope from 1 to 1.2 GHz if RF actives with GaAs hybrids are replaced with the new actives. Table 3 presents the results of analysis for possible level/slop increases while maintaining the levels with 1 GHz within ±1 dB of the current levels under the assumption that the output power capability of GaN hybrids is approximately 2 dB higher than that of GaAs hybrids for the same level of performance. For RF actives with older hybrid generations, the output power capability increase was larger so when analyzing the network for BW upgrades, the generation of the RF active hybrids should be determined and considered.

	Slope	Levels at 54 MHz	551 MHz	999 MHz	1191 MHz	Load	Composite Power	Technology
Bandwidth	[dB]	[dBmV]	[dBmV]	[dBmV]	[dBmV]		[dBm]	
54 to 1002 MHz	14.0	37.1	44.4	45.0	NA	79CW + 75QAM(-6 dB)	14.9	GaAs
54 to 1002 MHz	14.0	37.1	44.4	45.0	NA	79A + 75QAM(-6 dB)	13.4	GaAs
54 to 1002 MHz	14.0	31.1	38.4	45.0	NA	154QAM	13.0	GaAs
54 to 1002 MHz	14.0	37.1	38.4	45.0	NA	384 MHzP (<500 MHz) + 90QAM(-6 dB)	14.5	GaAs
54 to 1194 MHz	17.0	31.1	38.4	45.0	48	192 MHz + 154QAM	16.0	GaN
54 to 1194 MHz	17.0	31.1	38.4	45.0	48	384 MHzP (<500 MHz) + 122QAM(-6 dB)	16.8	GaN
54 to 1002 MHz	18.0	39.6	48.9	51.5	NA	79CW + 75QAM(-6 dB)	19.9	GaAs
54 to 1002 MHz	18.0	39.6	48.9	51.5	NA	79A + 75QAM(-6 dB)	18.8	GaAs
54 to 1002 MHz	18.0	33.6	42.9	51.5	NA	154QAM	18.5	GaAs
54 to 1002 MHz	18.0	39.6	42.9	51.5	NA	384 MHzP (<500 MHz) + 90QAM(-6 dB)	19.6	GaAs
54 to 1194 MHz	21.5	33.6	42.8	54.8	48	192 MHz + 154QAM	22.0	GaN
54 to 1194 MHz	21.5	39.1	42.3	54.3	48	384 MHzP (<500 MHz) + 122QAM(-6 dB)	22.0	GaN

Table 3: Slopes and Levels in RF Actives of the HFC Plant: GaAs versus GaN

In House Wiring and CPE

The last section of the HFC network, in-house wiring, is out of the control of the network operator. It introduces the largest variability of performance due to several factors:

- Where in the cascade of RF amplifiers the CPE is located (some CPEs may be connected to the first line extended off the node, the others at the end of the cascade);
- Drop length variability to the house entrance;
- The in-house wiring splitting ratio and length of in-house wiring;
- The quality of the in-house wiring components.

The data presented during the standardization effort of advanced PHY for HFC access technologies support the high performance variability picture in both the downstream and upstream. In the downstream, the median SNR levels are close to a 36 dB value ±1 dB. There are however some CPEs with performance lower than 33 dB SNR. This number is bound to increase with an extension in the forward bandwidth unless the operators make an effort to lower the variability and improve the performance.
This variability should be lower in FD (passive coax) HFC networks as the cascade origination point difference disappears in this architecture.

Drop length variability will always be there but except for extreme cases it can be accounted for in the design rules.

The two last items are the most difficult to address but with integration/consolidation of CPE devices, gradual elimination of analog cable TV channels and improvements in home networking technologies, the needs for many CPE devices connected directly to the HFC network via extensive splitting to serve a number of rooms with different services, currently requiring separate CPEs, will gradually diminish.

In the extreme cases of high in-house wiring loss, the operator has an option to deploy drop amplifiers, uni- or bidirectional, depending on needs.

The ultimate solution would be deployment of residential gateways with the signal terminating there. This would allow for elimination of 10 dB loss of in-house wiring.

This, however, becomes critical for significant bandwidth expansion networks. For the incremental capacity expansion summarized in Table 1, the existing CPE configuration will cause only limited problem in the networks that were optimized for QAM-only load. In the networks with large numbers of analog channels, operating today reliably for both analog and digital services, the levels received by analog and digital CPEs and consumer electronic devices allow for the bandwidth expansion presented in Table 1 if they can be supported by other segments of the network.

Recap of the Existing 1 GHz Network Capacity Limits

The results presented in Table 1 show that the existing network capacity can be increased incrementally but the expansion of the capacity depends on how the network was designed.

If all positives align:

- All analog channels can be eliminated in the network designed to carry them reliably and replaced in large part by a priority load of RF digital channels;
- The expansion of bandwidth in analog fiber links is possible;
- The RF actives in N+x (x>3) use technology pre-dating GaN hybrids;

- All legacy digital channels can be replaced by Advanced PHY channels;
- The existing frequency split between downstream and upstream is not changed;
- The in-house wiring supports reliably services carried on the existing analog and digital channels;

then the network capacity can be increased from the reference number of 6.2 Gbps to above 10 Gbps.

Some of the conditions listed above are related to the transition from the existing services to the digital services on Advanced PHY (analog channel reclamation, replacement of legacy digital channels with Advanced PHY digital channels) and some on the technology capability. The analog fiber links are one of the major contributors to the technology related condition. Their elimination would enable remedying the remaining technology obstacles to capacity expansion. Indeed, it would allow for significantly larger bandwidth expansion than the maximum possible in the existing HFC network (without a distributed forward architecture) even if all the conditions for this expansion are met. This would allow for significant capacity expansion without the need to replace the existing legacy-signal based services.

DISTRIBUTED ARCHITECTURES

Telecommunication Networks

In data networking systems, distributed architectures are ubiquitous. Examples include the Internet, cellular voice networks, and aircraft control systems. Benefits of the distributed architecture include greater reliability (no single point of system failure), better cost-efficiency (small clusters of functionality rather than an expensive monolithic system), and easier scalability and manageability.

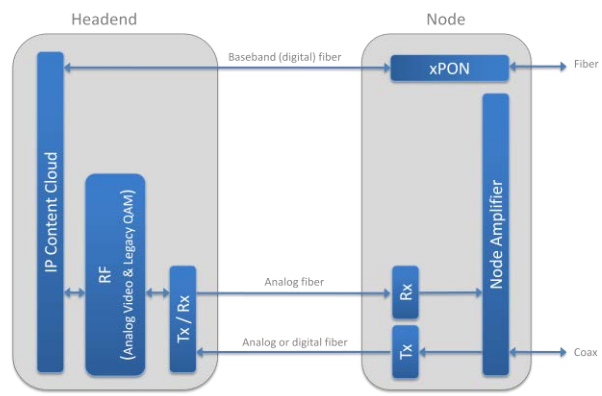

Figure 2: Headend-to-Node Block Diagram

Distributed Architecture Trends in HFC Networks

Until recently, the computational power required to process data and digital video and convert it to RF signals was so power- and space-intensive that it could only be done at the headend. A new generation of processors, computer memory, and FPGAs make it feasible now to migrate many functions further into the HFC network. The Distributed Broadband Access Architecture, in which downstream RF signals are generated (and upstream RF signals are terminated) remotely in the node, is an elegant concept that has long been discussed in theory, but data networking costs and harsh conditions in the node have made it impossible to implement until now. Space and power in the node are limited, and operating temperatures can easily exceed 70º C, but advances in silicon processes and capabilities, as well as data networking, make it possible now to implement certain RF-related functional blocks in a compact, low-power, temperature-hardened footprint. The most difficult and expensive functionalities to implement compactly are software-related—packet processing, filtering, and switching—or buffer and storage-related, but since these functionalities are well-suited to headend-based aggregation, leaving them at the cable headend for now not only simplifies implementation but also causes the least amount of operational change, and is also desirable from a software stability and maintenance perspective. Nonetheless, a few more generations of silicon will frame the option of an all-Ethernet HFC as merely a matter of preference, rather than a technological difficulty.

Happily, due to both the modular nature of modern fiber optic nodes, and also the frequency-based service channelization (a.k.a. frequency division duplex, or FDD), distributed digital functionality can be deployed as a simulcast or top-band frequency addition, in parallel with existing services. In addition, since downstream and upstream services are deployed on separate wavelengths, node-based modulation and demodulation can be introduced independently (see Figure 3).

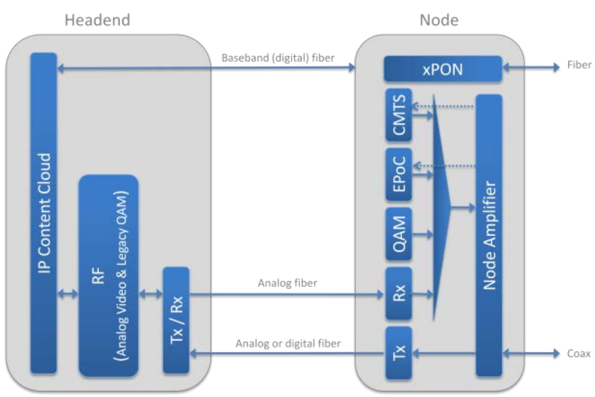

Figure 3: Migration of Services to the Node

As the analysis in the next section shows, the biggest benefit accrues from moving downstream RF generation from the headend to the node. Using today's chips and components, it is possible to fill the entire forward band, from 54 MHz all the way to 1,000+ MHz, with universal (broadband, narrowcast, and data) QAM-RF channels generated in the node. It is also possible with today's technology to place node-generated QAM channels all the way up to 1,800+ MHz. Future RF modulation standards and networking protocols can be added modularly (both physically in the node, and spectrally) as price-benefit tradeoffs dictate. For example, a Node QAM solution could be used to deliver legacy broadcast and narrowcast services in the 54 MHz – 1 GHz range, paired with a next-generation Ethernet-over-Coax solution above 1 GHz to deliver newer IP-based services. In certain cases, it may even be possible to directly convert Node QAM resources to other modulation protocols, although dense implementation of some of the more complex protocols currently under development will certainly require new silicon in order to meet power, space, and price constraints.

On a side note, the seemingly trivial question of what format to use for the digital transmission of bits between headend and node actually has cost and security implications. Options range from implementing a full multi-service TCP/IP router, to a Layer 3 or Layer 2 Ethernet switch, all the way down to a simple interleaved serial protocol. These solutions have a sliding scale of software complexity, introduce corresponding levels of delay and jitter, and require proportionate amounts of de-jittering buffer (which also consume space and power). Ultimately, the ability to connect the HFC plant directly to the headend IP network, or even the Internet, will be irresistible (see Figure 4). There will be certain costs, but many of those costs will become moot within a few generations of silicon, although vendors and operators should be mindful that the benefits of open Internet networking standards are paired with the responsibility to protect content from network-based theft, and the network itself from malicious intruders.

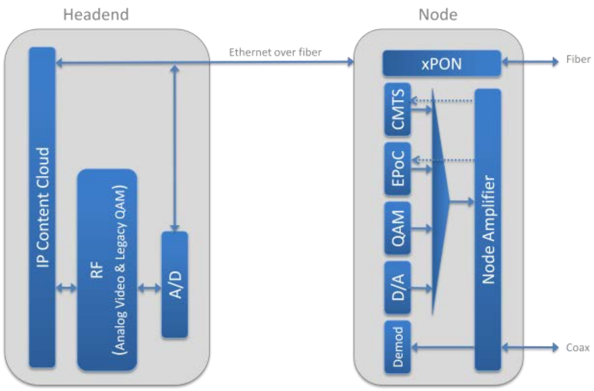

Figure 4: All-digital Broadband Access Architecture

In the upstream direction, some of the benefits of the Distributed Broadband Access Architecture can be realized simply by using digital return technology, already widely available from the top Optical Access Equipment vendors, and deployed in many hundreds of thousands of nodes. However, implementing upstream demodulators in the node have other benefits. For example, using a burst-receiver in the node to terminate the DOCSIS® return path not only reduces the noise funneling problem, but also simplifies the DOCSIS timing requirements and eliminates the need to keep cable modems within 100 miles of the headend. This approach meshes well with the goal of an all-Ethernet transport between headend and node.

Node Real-Estate

The Distributed Broadband Access Architecture is enabled through the precious real-estate owned by HFC operators: optical nodes. These robust and environmentally hardened enclosures with full remote monitoring and control provide significant advantage over CO/OTN-based topologies. They are much more robust and much closer to the final user and evolved into sophisticated yet reliable enclosures that are capable of supporting a multitude of Distributed Broadband Access Architectures in so called Virtual Hub configurations. Every optical node can be converted to a center of signal processing with all benefits of Distributed Broadband Access Architectures. The operational and technical and purely capacity expansion capability of this approach is unparalleled in any other telecommunications industry. Wi-Fi networks provide a similar approach but without the broad range of possible applications. Indeed, Wi-Fi networks can also benefit from and be supported by the optical node-based architecture.

CAPACITY EXPANSION WITH DISTRBUTED ARCHITECTURE

Bandwidth Expansion

Capacity expansions presented in Table 1 are mostly related to replacement of the existing legacy signals (analog and digital) with the Advanced PHY signals with their higher efficiency of bits/Hz. If there is a need to keep the legacy signals (and there will be for a while), the only viable solution to significantly increase the capacity is bandwidth expansion.

Bandwidth Efficiency

In any bandwidth, the material increase in capacity depends on improvement in network performance. As presented, the existing HFC networks for the performance they are designed achieve high efficiency already. This efficiency can be incrementally increased without improving the network performance by 20% to 30% with the introduction of Advanced PHY. However, the sizeable improvement in bandwidth efficiency will be gated by the capability of improving the network performance.

Both capacity expansion options are greatly advanced with distributed architectures.

DISTRIBUTED ARCHITECTURE IMPLEMENTATION CASES

Analysis Case Examples: Characteristics

Several examples of operating systems were selected for analysis. Characteristics of a few representative samples are collected in Table 4. These systems were re-designed with a distributed architecture for expanded bandwidth and their performance was modeled based on the design data. The list of the re-design cases analyzed is presented in Table 5.

Cascade	Existing High Fwd Freq [MHz]	MSO	System	Node	Plant Miles	HP	Density HP/MI
N + 5	750	A	Moscow	A149	11.91	758	64
				A168	10.55	955	91
				Total	22.46	1713	76
N + 5	860	B	Stalingrad	B15	7.94	556	70
				B52	6.81	628	92
				Total	14.75	1184	80
Fiber Deep (N + 0)	860	C	Kaliningrad	C004	1.10	92	84
				C005	1.26	116	92
				C006	0.95	50	53
				C007	1.02	123	121
				C008	0.90	64	71
				C009	0.89	77	87
				C010	0.92	64	70
				C011	1.24	120	97
				C012	1.01	111	110
				Total	9.29	817	88
Fiber Deep (N + 0)	1 GHz	D	Leningrad	D01	0.83	68	82
				D02	1.02	133	130
				D03	0.57	41	73
				D04	0.48	33	68
				D05	0.64	43	68
				D06	0.55	18	33
				D07	0.46	27	58
				D08	0.37	23	62
				D09	0.65	31	48
				D10	0.58	97	169
				D11	0.41	117	283
				D12	0.66	26	40
				D13	0.60	41	69
				D14	0.57	9	16
				D15	0.71	54	76
				D16	0.56	27	49
				Total	9.63	788	82

Table 4: Examples of Networks Analyzed and Re-Designed with Distributed Architecture

Original Design Freq & Type	C – New designs option 1 - with original house wiring								Number of Cases
	860	1002	1100	1200	1300	1400	1600	1800	
750 N+5	2	2							4
860 N+5		2	2	2					6
860 FD		9	9	9					27
1GHz FD				16	16	16			48
CPE	**Total number of test cases for individual network**								**85**
Original design Freq & Type	**G – New designs option 2 – without in-house wiring (with the gateway)**								
	860	1002	1100	1200	1300	1400	1600	1800	
750 N+x			2	2					4
860 N+x					2	2			4
860 FD					9	9	9		27
1GHz FD						16	16	16	48
Gateway	**Total number of test cases for individual network**								**83**
Grand-total number of test cases for individual network with different project specs									**168**

Table 5: Examples of Networks Analyzed and Re-Designed with Distributed Architecture

Downstream Performance of the RF Cascades

For apparent reasons, the RF cascade performance modeling for all extended bandwidth designs are presented for cases of the original 750 MHz N+5 and 860 MHz N+5 designs only.

The designs were performed with “Lode Data” design software under the following assumptions:

- No re-spacing of RF amplifiers;
- RF amplifiers with higher output level capability:
- 3 dB higher at the highest frequency for added bandwidth up to 1.2 GHz (assumes replacement of or modulation related de-loading from analog channels),
- 2 dB higher for bandwidth extension above 1.2 GHz (the analog channels are a smaller part of the total load and their removal or modulation effect is minimized),
- The amplifier gain is increased to support higher output levels at lower input levels defined by the design,
- This assumption may be too conservative for 750 MHz systems and for some 860 MHz systems with RF amplifiers dating back pre-GaAs period;
- Noise factor of RF amplifiers of any bandwidth is the same at the highest frequency as for 1 GHz amplifiers at the highest frequency;
- Passive loss above 1 GHz is the same as at 1 GHz. (The advanced passives designed for 2 GHz and above have usually lower losses above 1 GHz than the 1-GHz traditional passive loss at 1 GHz);
- Above 1.2 GHz, some traditional passives would exhibit excessive loss or suck-outs,
- The plate replacement is anticipated for the passives that do not meet the assumption;
- The input level is padded only if it exceeds the maximum required input level to the first hybrid and the remaining padding/gain alignment is performed interstage.

Table 6 summarizes the CNR modeling results for the cascades of RF amplifiers for the designs completed under these assumptions.

Design Scenario	Original Design Upper Frequency	Worst/Best Case CNR for Longest Cascade		Potential Capacity Increase (Legacy Signals/Adv. PHY – non-priority load)
		New Upper Frequency	Original Upper Frequency	
	[MHz]	[dB]	[dB]	[Gbps]
750 MHz N+5 "Moscow" system	750	49.8/52.2	49.8/52.2	0 (max original capacity: 6 Gbps with 384 Adv. PHY priority load)
	860	49.0/53.0	50.7/53.0	0.8/1.1
	1002	47.9/51.3	48.8/51.3	1.7/2.2
	1100	44.2/48.7	46.9/49.6	2.3/3.1
	1200	43.1/47.7	46.3/48.8	3.0/4.1
860 MHz N+5 "Stalingrad" System	860	56.0/56.2	56.0/56.2	0 (max original capacity: 6.8 Gbps with 384 Adv. PHY priority load)
	1002	55.9/56.1	56.6/56.9	0.9/1.2
	1100	55.2/55.7	55.3/55.9	1.5/2.0
	1200	54.1/54/1	54.5/54.6	2.2/2.9
	1300	52.5/53.0	54.2/54.6	2.9/3.8
	1400	51.8/51.9	53.5/53.6	3.5/4.6

Table 6: CNR Performance of Expanded Bandwidth RF Amplifier Cascades in Analyzed Networks

The results of the design and modeling show reasonable degradation in performance if the bandwidth extension is relatively low (up to 25% and 30% and increases at accelerated rate above these ranges). The impact of the decrease in performance will be modeled at EOL performance modeling.

<u>Drop Performance Modeling</u>

Table 7 summarizes the design and modeling results for the drop section of the HFC plant. The drop was assumed to be 150 feet of RG6 cable (close to the median value of the drop length based on analysis presented during Advanced PHY standardization activities) with a single coupler following the drop to the house and 50 feet of RG59 cable wiring inside the house.

Table 7 presents median input level designs. Based on the authors' experience, and industry studies, the median input levels to CPE devices in the downstream direction are close to –3 dBmV for digital RF channels. The results of modeling, except for FD designs that take advantage of the absence of RF amplifier CNR contribution, closely match this number. The CNR performance was calculated based on the assumption of 10 dB NF for residential gateway and CPE devices.

Original design Freq & Type	New Design Frequency	Median CPE Input Levels at: Highest Freq./Original Design	Median CPE Input Levels at: New Design /Highest Original Frequency	Median CPE Input Levels at: Highest Frequency of New Design	CPE CNR at Median Input Levels at: Highest Freq./Original Design	CPE CNR at Median Input Levels at: New Design /Highest Original Frequency	CPE CNR at Median Input Levels at: Highest Frequency of New Design
	[MHz]	[dBmV]	[dBmV]	[dBmV]	[dB]	[dB]	[dB]
750 N+5	860	-3.02	-2.12	-2.63	44.22	45.12	44.61
	1,002	-3.02	-4.12	-5.15	44.22	43.12	42.09
860 N+5	1,002	-2.22	-1.42	-2.26	45.02	45.82	44.98
	1,100	-2.22	-2.62	-2.12	45.02	44.62	45.12
	1,200	-2.22	-4.42	-3.92	45.02	42.82	43.32
860 FD	1,002	-6.39	-6.09	-7.65	40.85	41.15	39.59
	1,100	-6.39	-7.59	-9.37	40.85	39.65	37.87
	1,200	-6.39	-9.59	-11.56	40.85	37.65	35.68
1GHz FD	1,200	-2.70	-3.71	-2.16	44.54	43.53	45.08
	1,300	-2.70	-5.24	-3.34	44.54	42.00	43.90
	1,400	-2.70	-7.18	-6.59	44.54	40.06	40.65
		Median RG Input Levels at			RG CNR at Median Input Levels at		
750 N+5	1,100	2.68	-0.23	-0.20	49.91	47.01	47.04
	1,200	2.68	-1.03	-1.50	49.91	46.21	45.74
860 N+5	1,300	4.04	1.14	1.39	51.28	48.38	48.63
	1,400	4.04	0.44	1.90	51.28	47.68	49.14
860 FD	1,300	-0.13	-4.23	-5.18	47.10	43.01	42.06
	1,400	-0.13	-5.52	-6.34	47.10	41.72	40.90
	1,600	-0.13	-6.43		47.10	40.81	
1GHz FD	1,400	4.42	-0.12	2.30	51.65	47.12	49.54
	1,600	4.42	-0.56	0.52	51.65	46.67	47.76
	1,800	4.42	-1.65	-1.14	51.65	45.59	46.10

Table 7: Drop Statistics and CNR Performance for Analyzed Cases.

		With Re-Designed Analog Fiber Link for Better Performance			Ditsributed Architecture Optimized for 1024-QAM with LDPC			Ditsributed Architecture Optimized for 4096-QAM with LDPC		
Original design Freq & Type	New Design Frequency	End-To-End Downstream Performance: CPE with In-House Wiring			End-To-End Downstream Performance: CPE with In-House Wiring			End-To-End Downstream Performance: CPE with In-House Wiring		
		Highest Freq./Original Design	New Design /Highest Original Frequency	Highest Frequency of New Design	Highest Freq./Original Design	New Design /Highest Original Frequency	Highest Frequency of New Design	Highest Freq./Original Design	New Design /Highest Original Frequency	Highest Frequency of New Design
	[MHz]	[dBmV]	[dBmV]	[dBmV]	[dBmV]	[dBmV]	[dBmV]	[dBmV]	[dBmV]	[dBmV]
750 N+5	860	36.81	37.00	36.83	40.07	40.49	40.12	42.16	42.85	42.23
	1,002	36.81	36.54	36.23	40.07	39.51	38.93	42.16	41.28	40.43
860 N+5	1,002	37.12	37.25	37.11	40.75	41.05	40.74	43.32	43.88	43.29
	1,100	37.12	37.04	37.12	40.75	40.58	40.76	43.32	43.00	43.34
	1,200	37.12	36.66	36.76	40.75	39.75	39.97	43.32	41.65	42.00
860 FD	1,002	36.18	36.28	35.71	38.78	38.97	37.96	40.23	40.49	39.12
	1,100	36.18	35.74	34.92	38.78	38.00	36.70	40.23	39.17	37.54
	1,200	36.18	34.81	33.67	38.78	36.54	34.94	40.23	37.34	35.48
1GHz FD	1,200	37.13	36.93	37.22	40.69	40.25	40.90	43.21	42.45	43.60
	1,300	37.13	36.55	37.01	40.69	39.46	40.42	43.21	41.21	42.73
	1,400	37.13	35.90	36.11	40.69	38.28	38.66	43.21	39.54	40.06
		End-To-End Downstream Performance: RG			End-To-End Downstream Performance: RG			End-To-End Downstream Performance: RG		
750 N+5	1,100	37.43	36.98	36.62	41.50	40.44	39.67	44.78	42.77	41.52
	1,200	37.43	36.83	36.25	41.50	40.11	38.96	44.78	42.22	40.49
860 N+5	1,300	37.69	37.48	37.46	42.21	41.65	41.58	46.47	45.10	44.96
	1,400	37.69	37.41	37.47	42.21	41.45	41.62	46.47	44.67	45.03
860 FD	1,300	37.50	36.81	36.56	41.57	39.99	39.49	44.94	42.03	41.26
	1,400	37.50	36.46	36.20	41.57	39.30	38.81	44.94	40.97	40.27
	1,600	37.50	36.17		41.57	38.76		44.94	40.20	
1GHz FD	1,400	37.82	37.50	37.71	42.44	41.58	42.13	47.12	44.95	46.25
	1,600	37.82	37.45	37.56	42.44	41.45	41.75	47.12	44.67	45.33
	1,800	37.82	37.30	37.37	42.44	41.10	41.27	47.12	43.96	44.30

Table 8: End-To-End Downstream Performance

Downstream End-To-End Performance

The results in Table 8 for the EOL (a.k.a. end-to-end) performance with analog fiber links show higher than expected results but this is mostly due to the many assumptions on the analog fiber link used during the analog fiber link capacity analysis presented in Table 1 (including the assumption that the links were designed to carry high numbers of analog channels). The test results also indicated that there is very little change in the performance from improving drop installation performance (moving CPE to RG). This indicates that the analog fiber link performance is the limiting factor and the critical component of the network. The assumptions used can be possibly met in a single wavelength 1310 nm analog links but very difficult to meet in multiwavelength applications.

On the other hand, a distributed architecture allows extended bandwidth with performance (median numbers) capable of supporting 4096-QAM at the highest frequencies to and in several cases beyond the limits analyzed even if the performance of the PHY output is at the level optimal for 256-QAM or 1024-QAM with LDPC (the implementations with guaranteed 43 dB SNR/MER and typically >44 dB SNR/MER).

Without relocation of the CPE devices into RG, 750 MHZ N+5 network bandwidth can be extended to 1002 MHz and entire bandwidth can support 4096-QAM (median). After relocation of the CPE devices into RG, the bandwidth can be extended to 1.2 GHz and the network could support 4096-QAM in the most part of that bandwidth. The analysis results also show that under the same two drop topologies, the 1GHz FD designs can be extended in BW to 1.4 GHz and 1.8 GHz respectively while supporting 4096-QAM with LDPC within most of those bandwidth ranges.

With the distributed architecture remote PHY optimized for higher performance (49 dB MER/SNR), these bandwidths can be easily extended or the network could provide high operational margins and support excessive loss drops

Note that in 1.8 GHz network, the downstream capacity (after replacing all legacy signals with Advanced PHY signals) can exceed 17 Gbps (starting from 168 MHz up to 1.8 GHz). The results show that the downstream bandwidth can be indeed extended to 2 GHz.

Upstream Network Capacity Analysis

Two different approaches are discussed for the upstream capacity expansion:

1. Moving downstream/upstream split band to upper frequencies (low split band),
2. Placing the new upstream bandwidth over the top frequency of the downstream bandwidth (over the top split band).

The pros and cons have been broadly discussed by the industry. The few (major) cons are listed below:

1. For split band relocation:
    a. Sizeable shift in upstream frequency would require remedies (also broadly discussed) up to and including legacy set top boxes with limited downstream OOB signaling agility;
    b. Alternative remedies range from placing simple frequency converters to complete OOB receivers attached to the set top boxes;
2. For over the top:
    a. High cable loss for the signal generated by CPEs, even if placed in residential gateways;
    b. Restriction to future downstream bandwidth expansion that would require a visit to the RF actives.

Both solutions could be implemented in TDD or FDD configurations. For over the top, TDD implementation addresses future forward capacity expansion (adding forward capacity by adding more or higher throughput TDD channels). The advantage of the FD architecture is that it can readily support both or either.

The advantages of the TDD approach will not be discussed in this paper. They are numerous and very well understood.

Unfortunately, today native HFC technologies do not support TDD. This is potentially a missed opportunity.

TDD technology, by implementing or maintaining some available capacity asymmetry, can be used over the top with upstream signals at much lower modulation levels and thus requiring lower power and lower performance. For example, using the same BW downstream and upstream but with 4096-QAM downstream and 64-QAM upstream would support asymmetric available capacity of 2-to-1 (if TDD is split equally between downstream and upstream but it is indeed flexible in how the capacity is used) but the upstream signal with the same FEC power strength as the FEC for the downstream signal could occupy an 18 dB lower SNR environment that can be implemented entirely or partially as lower level transmitting requirements in CPEs.

At lower frequencies with a traditional split band (high) approach, the lower part of the HFC bandwidth can be used for FDD and TDD (in FD networks). In this case, the lower loss would allow for fully symmetrical capacity availability in downstream and upstream without burdening CPE devices with extremely high output power requirements.

SUMMARY: BENEFITS OF DISTRIBUTED BROADBAND ARCHITECTURE

There are many operational and financial benefits to carrying only baseband digital signals in the fiber portion of the HFC network, rather than modulated RF signals. These include being able to leverage high-volume data networking equipment; reducing headend space and power consumption; increasing fiber reach; decoupling service group allocations from hard-wired RF combining networks; and eliminating the need for plant rebalancing, with "set and forget" channel power settings.

The most significant benefit, however, is that eliminating RF payload from the fiber unlocks as much as a three-fold improvement in the potential bandwidth capacity of the existing plant infrastructure. The plant modeling analysis above, using real design examples, shows that without re-spacing RF actives in a Fiber Deep and HFC network up to N+5 deployments, forward path capacity can reach nearly 20 Gpbs, given the aid of LDPC FEC encoding at the headend and more powerful GaN hybrid amplifiers in the node.

The source of this benefit is twofold. Firstly, by moving RF modulation and/or demodulation to the node, the entire "headend-quality" RF signal budget is available at the last mile, without incurring any of the traditional headend-side or optical transmission-related impairments. Secondly, there are system-wide benefits even during the transition period (while a node receives both digital optical signals and legacy analog optical signals) on top of bandwidth expansion. This is due to the fact that as video and data QAMs are migrated to the node, the decreased RF burden on headend optical transmitters enables increased OMI, which results in improved SNR at the node for legacy headend signals as well.

ABBREVIATIONS AND ACRONYMS

AGC	Automatic Gain Control
ALC	Automatic Level Control
AML	Amplitude Modulated Link
BC/NC	Broadcast/Narrowcast
BER	Bit Error Rate
BW	Bandwidth
CM	Cable Modem
CMTS	Cable Modem Termination System
CNR	Carrier-to-Noise Ratio
CO	Central Office
CPE	Customer Premises Equipment
CSO	Composite Second Order
CTB	Composite Triple Beat
dB	decibel
DOCSIS®	Data over Cable Service Interface Specification
E-O-E	Electrical-Optical-Electrical
EOL	End-of-Line
FD	Fiber Deep
FDD	Frequency Division Duplex
FEC	Forward Error Correction
FML	Frequency Modulated Link
FPGA	Field-Programmable Gate Array
GaAs	Gallium Arsenide
GaN	Gallium Nitride
Gbps	Gigabits per second
HFC	Hybrid Fiber Coaxial
IIN	Interferometric Intensity Noise
IP	Internet Protocol
LDPC	Low Density Parity Check
Mbps	Megabits per second
NA	North America
NF	Noise Factor
NTSC	National Television System Committee
OFDM	Orthogonal Frequency-Division Multiplexing
OMI	Optical Modulation Index
OOB	Out-of-Band
OTN	Optical Terminal Node
PAPR	Peak-to-Average Power Ratio
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RFoG	Radio Frequency over Glass
RG	Residential Gateway
RIN	Relative Intensity Noise
RMS	Root-Mean-Square
SNR	Signal-to-Noise Ratio
TDD	Time Division Duplex
TX	Transmitter
VSB	Vestigial Sideband

ENHANCING DOCSIS SERVICES THROUGH NETWORK FUNCTIONS VIRTUALIZATION

Andrew Smith, Colby Barth
Juniper Networks
Saif Rahman
Comcast

Abstract

The DOCSIS system does not directly define a service offered to a subscriber. The use of classifiers, service-flows and the like define the parameters around a service – typically IP high-speed-data, L3VPN, or L2VPN but does not make any attempt to define, or enhance, the actual service delivered over DOCSIS. This paper will examine the role that Network Functions Virtualization (NFV), an effort within the Software Defined Networking (SDN) community, can play as an enhancement to data services delivered via DOCSIS. We will examine benefits of NFV, propose an integration scheme of NFV and DOCSIS, and explore some potential use cases of this enhanced DOCSIS services architecture. This integration can increase network efficiency, enable new revenue opportunities, and permit rapid service innovation.

NFV OVERVIEW

Network Functions Virtualization is an effort originally proposed by the European Telecommunications Standard Institute (ETSI), a group of service providers and operators at the SDN and OpenFlow World Congress in 2012. [1]. The objective of NFV is to leverage modern generic computation and software virtualization systems to enable a new platform upon which to build high-touch network services and revenue models. The number of proprietary hardware appliances is reduced, with an increase of standard x86-style compute platforms and software implemented on data and control plane packet processing.

The benefits of NFV are, but are not limited to:

- Reduction in equipment costs and power through economies of scale
- Increased innovation velocity through software development
- Elastic scaling of a services-plane within a network
- Open and robust services innovation

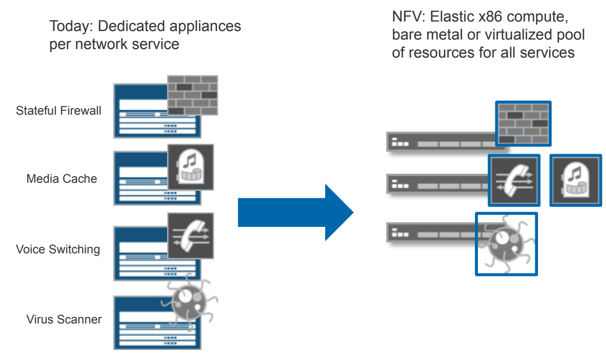

BENEFITS OF DOCSIS WITH NFV

Network Functions Virtualization presents an innovation opportunity for cable operators. Previous models of network service architecture required large, centralized service data centers (i.e. voice switching) or highly distributed, purpose built appliances (i.e. video caching), both with economic and operational challenges. To maintain large central data centers in cable environments requires service rates and take-up percentages to justify the build and operating costs. Distributed appliances have similar challenges, in that they must offset enough capex from "business as usual" – simply deploying more bandwidth – to become feasible.

The virtualized aspect of NFV permits one investment, in generic x86 computing resources, to be leveraged enabling many services. Where previous network services were instantiated on dedicated, standalone appliances – each with their own vendor, support, operational model and depreciation schedule – now one investment can be made to support all services. Services themselves run as virtual machines in an x86 compute environment. Individual services can scale just as in a modern data center, through the allocation of virtualized memory, CPU, storage and network policy, as well as horizontally by adding additional load-balanced virtual machines.

The generic nature of this new services infrastructure enables us to propose an integration of DOCSIS and NFV, such that the cable operator can enable services on a per-service flow, per-cable modem or per-VPN nature. These services would be provisioned using the standard DOCSIS tools and existing TFTP and DHCP mechanisms of standard cable modem registration.

NFV AND DOCSIS INTEGRATION

The figure illustrates a potential deployment scenario in a typical cable hub site with local and remote x86 service elements. The "ER" in the figure is the CCAP "Edge Router." In an Non-Routing (NR) CCAP architecture, the ER manages IP sessions, IP routing, IP/MPLS servces and implements NFV, while the CCAP handles specific MAC and PHY interfaces for DOCSIS. In an Integrated-CCAP or traditional CMTS architecture, all of this hardware and software functionality must be built in to a single chassis. Either architecture is capable of integrating NFV with DOCSIS. The differences between NR-CCAP and Integrated-CMTS/CCAP architectures are beyond the scope of this paper.

In this example, three services have been virtualized: Analytics and Lawful

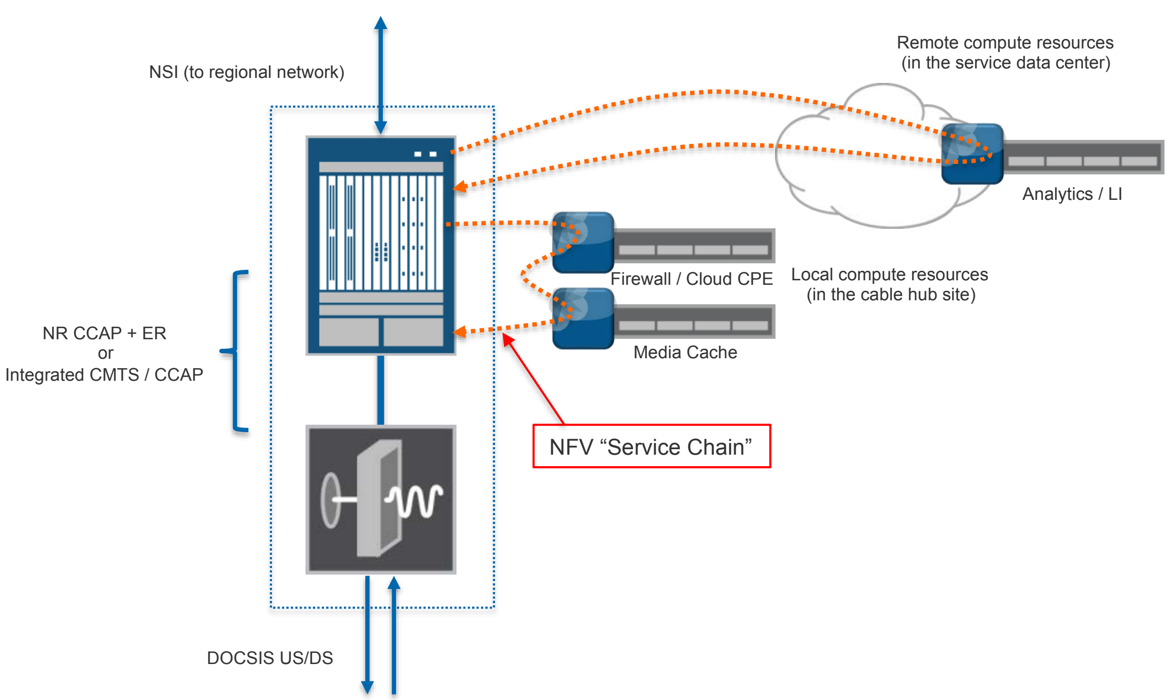

Intercept, Firewall and Cloud CPE, and a Media Cache. The Lawful Intercept function has been virtualized in a centralized manner, the other services are local to the hub site. The architecture permits this flexibility – some services may have a very low usage rate where economies of scale will find the optimal placement to be in a central location. Policy may also dictate a centralized location, such as data security or integrity.

Services that are part of the data plane and potentially latency sensitive should be placed in the cable hub site. This is for performance optimization as well as to restrict the “tromboning” effect of traffic flows to local, short reach interfaces and avoid metro or long-haul optical networks.

All of these services are implemented on standard x86 computing environments. As the network grows and changes, simply rearranging software elements can shift where these services are implemented. No service is permanently fixed to a particular location or resource in the network. The operator is able to adapt this system to their particular use case, load or traffic pattern, or policy requirements.

The combination of services enabled for a flow of data is called a “service chain.” The service chain may contain one or more elements to perform actions upon the data flow, and they can be positioned local to the hub site or remotely across the IP/MPLS metro network.

A flow of traffic may be routed through a Lawful Intercept and Media Cache. Another flow could go through only the Cloud CPE. The construct of the service chains are entirely arbitrary and up to the discretion of the cable operator. The x86 compute devices and their resident NFV service functions are connected to the network using standard Ethernet and/or MPLS VPN implementations, depending on the proximity of the services defined in the service chain to the Edge Router.

The particular services deployed in a service chain is up to the operator. A typical cable operator may pre-define several service chains for given service levels to subscribers.

Chain	Flow	Description
Set A	Internet > Subscriber	Today’s status quo, HSD service
Set B	Internet > Cache > Subscriber	An HSD subscriber with video cache selectively enabled before the downstream
Set C	Internet > LI > Cache > Subscriber	HSD subscriber, data shipped off to remote LI collection before being routed to local cache then on to the downstream
Set D	Internet > Cloud CPE > Cache > Subscriber	Subscriber traffic bounces through two service elements before being routed downstream

PROVISIONING NFV AND DOCSIS

Service sets will be pre-defined once in the configuration of the CMTS/CCAP or the ER of an NR-CCAP system. These service sets need only be configured once. One instance of a service-set can be applied to multiple simultaneous DOCSIS service flows.

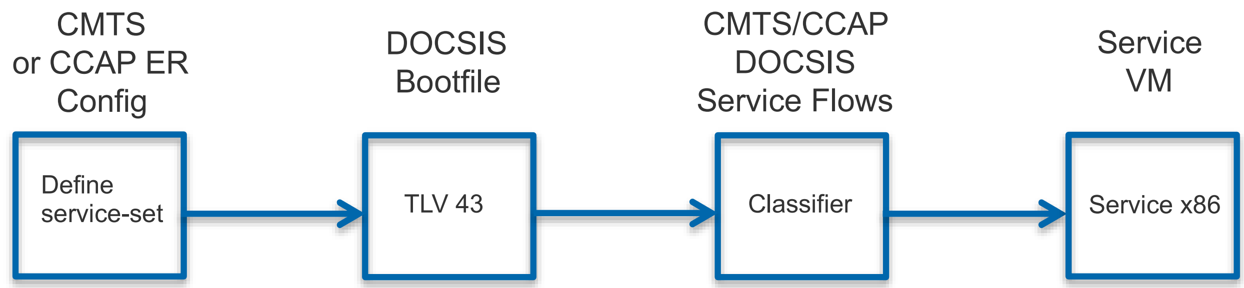

This paper proposes a new Type-length-value (TLV) in the DOCSIS boot file to associate a specific DOCSIS service-flow with a specific NFV service chain. We initially propose a VendorSpecific TLV but this is envisioned to become part of an enhanced, standards based DOCSIS service definition mechanism.

Once in place, when the cable modem registers to the CMTS or NR CCAP, the specific service-flows defined in the DOCSIS boot file will be associated with NFV service-sets, and the data classified in those service-flows will be subject to the actions defined in the service-set.

The use of DOCSIS Service Class Names could also be used to abstract the NFV portion of the ER confirguration from the cable modem boot file's specific DOCSIS service-flows, depending on operator implementation and potential operator billing preferences. This would also permit more granular definition of the DOCSIS to NFV relationship in the running config of the ER or Integrated system. [2]

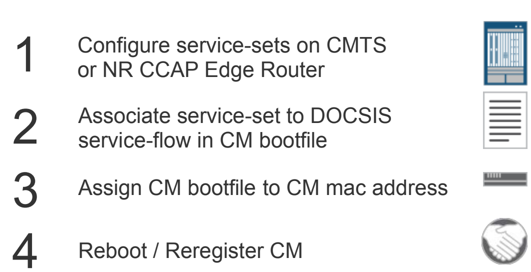

Future architectures may be able to dynamically associate DOCSIS service flows to NFV service-sets without rebooting the cable modem, through the use of SNMP, PCMM or similar centralized signaling mechanisms.

The use of a VendorSpecific TLV ensures that if a CM registers to a legacy CMTS that does not have NFV support that the TLV is subsequently ignored. The CM will be allowed to come online, but there will be no benefits of the NFV infrastructure. [3]

SERVICE CHAIN CONFIGURATION

It is envisioned that the service-chains are configured universally across a cable operator's network to permit cable modem mobility, i.e. a cable modem is not tied to a particular CMTS or hub site. The service chains themselves need not operate identically – in one hub site an LI function could be local, in another it could be remote, but from the perspective of DOCSIS the operation is identical.

The NR-CCAP ER or CMTS will know the service-set, and location of the service instances, based on the running config of the system. A system in a smaller hub site may be configured to tunnel certain NFV flows back to a central location, while a system in a large hub site may locate services locally. Where services are instansiated depends upon the network topology, service requirements and other critera defined by the operator.

The service-chain “Set C” from the previous examples, represented on a NR-CCAP ER or an integrated CMTS/CCAP, could be configured as follows. This is a representative XML schema:

```
<rpc-reply
xmlns:junos="http://xml.juniper.net/junos/13.3
X-cable/junos">
  <version>13.3X-cable [slt-builder]</version>
    <services>
      <service-set Set-C>
        <docsis-flow-id 101>
          <chain>
            <service-instance>
             <target>LI-DataCenter-HQ</target>
            </service-instance>
            <service-instance>
             <target>Cache-HTTP-Local</target>
            </service-instance>
            <service-instance>
              <end/>
            </service-instance>
           </chain>
         </docsis-flow-id>
       </service-set>
      <service-target>
        <name>LI-DataCenter-HQ</name>
         <location>remote</location>
          <access-via>
            <pwe3>
              <remote>
               <host>69.252.101.5</host>
                <vc-id>500</vc-id>
              </remote>
            </pwe3>
          </access-via>
      </service-target>
      <service-target>
       <name>Cache-HTTP-Local</name>
          <location>local</location>
           <access-via>
            <ethernet>
              <local>
                <interface>xe-5/3/2</host>
                <vlan-id>6</vlan-id>
              </local>
            </ethernet>
          </access-via>
       </service-target>
      </services>
```

This configuration defines:

- The DOCSIS service flow reference, in this case, the value 101. This value is also present in the DOCSIS boot file, and is the linkage between DOCSIS and the NFV chain,
- The service-chain, through a series of service-instance statements,
- The target elements of the service chain, one remote (LI-DataCenter-HQ) and one local (Cache-HTTP-Local) and the respective access methods to each chain element.

DOCSIS CONFIGURATION

The following is an example portion from a DOCSIS boot modem configuration.

```
DsPacketClass
{
      ClassifierRef 101;
      ServiceFlowRef 101;
      ActivationState 1;
      IpPacketClassifier
      {
            IpProto 6;
            DstPortStart 80;
            DstPortEnd 80;
      }
      GenericTLV TlvCode 43 TlvLength
8 TlvValue 0x0803ffffff010100

}
DsServiceFlow
{
      DsServiceFlowRef 101;
      QosParamSetType 7;
      TrafficPriority 4;
      MaxRateSustained 5000000;
}
```

This configuration file uses a DOCSIS Downstream Packet Classifier to select all TCP traffic on port 80 and push it to NFV service chain 101. This traffic is associated with Service Flow 101, and rate limited to 50 Mb/s.

The operator can choose any combination of service flows, rate limits, TOS markings or packet classifiers when choosing which traffic will, or will not, be subject to NFV functionality.

CONCLUSION

Network Functions Virtuzliation represents a potential new area of innovation in cable high speed data services. Service chains can be constructed by the operator to increase network efficiency, enable new revenue opportunities, or any combination thereof. The concepts presented in this paper illustrate a linkage between DOCSIS and the SDN community through software and services innovation. As the NFV and SDN communities mature, cable operators will be ready to take advantage of a next generation, revenue enhancing services-plane by following the concepts in this paper.

REFERENCES

[1] Chiosi, Margaret, et. al, (2012). *Network Functions Virtualisation, An Introduction, Benefits, Enablers, Challenges & Call for Action*, SDN and OpenFlow World Congress, Darmstadt-Germany

[2] Cablelabs, et. al, (2003). *Data-Over-Cable Service Interface Specifications, DOCSIS 1.1, Operations Support System Interface Specification, SP-OSSI-v1.1*, Cable Television Laboratories, Inc, Loiusville, CO.

[3] Cablelabs, et. al, (2007). *Data-Over-Cable Service Interface Specifications, DOCSIS 3.0, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0*, Cable Television Laboratories, Inc, Loiusville, CO.

CONTRIBUTORS

Stuart Mackie and Ken Gray, of Juniper Networks, contributed to this paper.

EXAMINING THE FUTURE EVOLUTION OF THE ACCESS NETWORK

Michael J. Emmendorfer and Tom Cloonan
ARRIS

Abstract

The MSOs are facing a decade of unprecedented change, driven by competition and consumer demand, which will transform the cable network end-to-end. This paper will focus on the evolution of the cable access network through this decade and beyond. We will examine the evolution of Coax to the Home (CTTH) and Fiber to the Node (FTTN) networking technologies and architectures that will emerge this decade.

The MSOs are maximizing capacity and performance of their current networks. However, MSOs are finding the improvements to operations can't yield any more bits per second per hertz (b/s/Hz). This paper will examine the future evolution of the access network, to identify some of the problems and proposed solutions to maximize b/s/Hz and "widen the pipe" to increase the capacity of existing and future spectrum.

Question #1: *Are cable networks "limited by" the RF video and data technologies, which are based on ITU-T J.83 annex A/B/C for downstream & CableLabs DOCSIS 2.0 Upstream?*

Question #2: *Are Cable Networks "limited by" the FTTN Optical Technology, which is based primarily on Amplitude Modulation (AM) optical technology to/from the node?*

Question #3 *Moving from AM Optics to Digital Optics for FTTN will force us to place PHY or MAC/PHY Access Layer Functions in the Node. What stays in the headend and what moves to the node? The industry will need to define a new access network architecture supporting digital connections between headend and fiber node. This new access network architecture will redefine the CCAP architecture and other headend platforms (e.g. Digital Optical Platforms) as well as the node platforms.*

INTRODUCTION

The capacity of the cable access network depends on several factors. These factors may include network operations, network architecture, spectrum selection, spectrum allocation, spectral load, RF technology and optical technology. We are finding that overall MSO operations and design practices will not be the limiting factor to maximize capacity. This paper suggests that today's cable network capacity or b/s/Hz is limited by the Radio Frequency (RF) Technologies supporting Digital Video and Data Services. This paper also suggest that as improvements are made to the RF Data Technology, such as DOCSIS 3.1, that the next limiting factor will be the Optical Technology to/from the HFC node. As our industry expands spectrum in the downstream and upstream the current optical technology will increasing become a limiting factor to maximize b/s/Hz. Additionally, as MSOs have a desire to reduce facilities and expand the optical distance between headend and fiber node this will also limit the system b/s/Hz, based on today's optical technology, Amplitude Modulation.

This paper will focus on the three core areas of the cable access network that are increasing becoming an integral part for maximizing spectral capacity. The sections that follow will be organized by the identified questions as stated above. In those sections we will provide additional detail to the scope of the problems and suggest solutions.

QUESTION #1: ARE CABLE NETWORKS "LIMITED BY" THE RF VIDEO AND DATA TECHNOLOGIES?

In order to determine the limiting factors of today's cable network an assessment of the RF technology attributes must be measured against the performance measurements of the cable network.

Overview of the "Current" RF Video and Data Technologies

The digital video and DOCSIS services deployed by cable operators around the world use an RF technology defined in Recommendation ITU-T J.83 and the four Annexes (Annexes A, B, C, and D). This standard defined the physical (PHY) layer technology used for digital video MPEG-TS and DOCSIS downstream specifications through version 3.0.

The main differences in the ITU-T J.83 annexes will be the channel coding and modulation specified, as well as the channel width. The highest order modulation in all versions is 256 QAM. A key attribute of the annexes is the selection of error correction technology. Annex A/C/D define a single error correction technology called Reed-Solomon. The ITU-T J.83 Annex B uses outer FEC called Reed-Solomon (R-S) and an inner FEC called Trellis Coded Modulation (TCM). The use of trellis coding in J.83 annex B is embedded in the modulation process. This means that J.83 annex B is more robust than the annex A/C/D versions. The impact of these differences in FEC means that J.83 Annex A/C will require about 2 dB better system performance than J.83 Annex B to support the same modulation format and assuming about the same code rate for each [2]. The applications that use J.83 annexes and the SNR (dB) requirements are found in tables 1 and 2 respectively.

J.83 Error Correction Technology [1]:

- ITU-T J.83 Annex A/C uses Reed-Solomon Downstream
- ITU-T J.83 Annex B uses Trellis Code Modulation (TCM) inner FEC and Reed Solomon (outer FEC)

ITU-T J.83-A	Euro-DOCSIS Annex A DVB-C
ITU-T J.83-B	DOCSIS Annex B Japanese DOCSIS Annex C ATSC/SCTE
ITU-T J.83-C	Japanese Digital Video

Table 1: ITU-T J.83 Applications

J.83 Annex	Coded Rate Assuming AWGN	Minimum Operating Recommendation
J.83-A	29 dB	32 dB
J.83-B	27 dB	30 dB
J.83-C	29 dB	32 dB
Assumptions: The coded value assumes a ~ 90% code rate		

Table 2: SNR (dB) for 256 QAM

The upstream RF data technologies are based on CableLabs DOCSIS 2.0 standard called A-TDMA and S-CDMA. These have different modulation and error correction technologies defined below:

Method	Error Correction Technology
A-TDMA	Reed-Solomon (R-S) 64 QAM
S-CDMA	Trellis Code Modulation (TCM) & Reed-Solomon (R-S) 128 QAM

Table 3: Upstream DOCSIS Error Correction Technologies

The use of Single Carrier with A-TDMA and Reed-Solomon with 86% Code will require at the slicer 22 dB, in words the CMTS upstream port [2]. If we add 7 dB above the slicer the systems requirements reach 29 dB. In our models used for the upstream we allocate 10 dB of margin about the slicer, thus 32 dB for A-TDMA 64 QAM. In practice, MSOs may target between 30 dB to 33 dB for a minimum operating recommendation.

A very important take-away from this section are the minimum operating recommendations. The downstream RF technology to operate 256 QAM for EuroDOCSIS assuming 92% code rate is 32 dB, DOCSIS 256 QAM code rate of 90.5% is 30 dB, and A-TDMA DOCSIS upstream with 64 QAM and an 86% code rate is 32 dB. These values represent the minimum operating recommendations for the cable operators to enable the highest order modulations possible with the current RF technologies. If the operator's network exceeds these minimum operating recommendations thresholds by greater than 3 dB, then the RF Technology is the limiting factor.

Examine the "Current" Cable Network Downstream Performance

This paper will determine if today's cable network or if today's RF data technology is the limiting factor for MSOs to achieve more b/s/Hz, in other words capacity.

In order to determine which is the limiting factor, the cable network or the RF technology, we need to understand the measurements of the cable network and the requirements of the RF data technology. The section above determined the minimum operating recommendations measured in dB to operate reliable service over time.

The author has received a contribution from Comcast Cable so that we may effectively assess the RF data technology against real-world network data. Dave Urban of Comcast Cable made this contribution and we thank him for this critical information for our study. Mr. Urban has completed pioneering research in measuring the performance of the cable network. His research as illustrated in figure 1, measures the downstream performance of 20 million cable modems. Mr. Urban has plotted these 20 million cable modems in a histogram by downstream signal to noise ration S/N in dB. His grounding breaking findings proved several key points to the cable industry:

1. The existing cable network supported the highest order modulation possible using DOCSIS / J.83B for all users.
2. Though the distribution of cable modem performance is vastly different, nearly all devices could support higher order modulation formats or more b/s/Hz if available.
3. This work is credited with convincing the industry to support in the future the use of Multiple Modulation Profiles (MMP). The use of MMP enables groups of modems sharing common SNR the ability to use the highest order modulation possible, maximizing b/s/Hz.

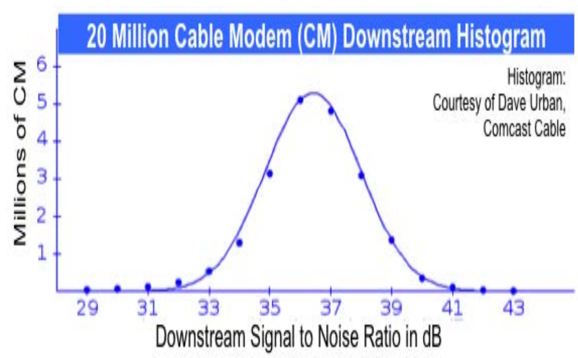

Figure 1: 20 Million Cable Modem Downstream Histogram

Estimating the Cable Network Upstream Performance

To estimate the use of the upstream cable plant and future spectrum splits ARRIS built a return path model. The ARRIS Upstream HFC Performance Model is an assessment of the Noise and Attenuation in the Optical and Coaxial Segments. The model considers many spectrum splits from 5-42, 5-85, 5-238, 5-500, and several Top-split spectrum options. The model proved that Top-split, placing the upstream above 900 MHz or much higher, was too costly and consideration for Top-split was abandon by the industry in late 2011. The ARRIS model has been vetted by MSOs, fellow suppliers, and was contributed to CableLabs.

The main purpose of the ARRIS Upstream HFC Performance Model is an analysis of the HFC Optical and Coaxial segment of the network under "normal operating conditions". In a given spectrum split the model estimates the system carrier to noise C/N to determine the highest upstream modulation type that may be used. The estimated C/N is then matched using OFDMA with LDPC and BCH error correction technology, planned for DOCSIS 3.1, and the highest modulation format given the assumptions used in the model. The figure below illustrates the areas of study in the model.

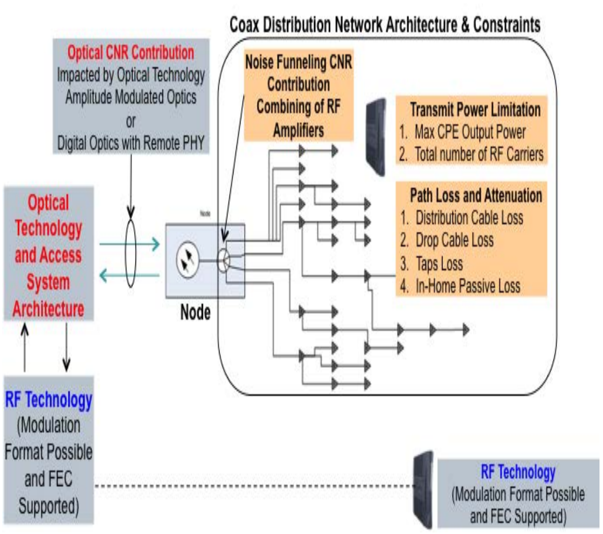

Figure 2: Major Considerations for Coaxial Network Performance

The key output of the model is the estimated system C/N and the modulation type as seen in the highlighted red boxes in Table 4. The model shows that the upstream could support higher order modulation to increase the b/s/Hz and overall system capacity if DOCSIS 2.0 had defined support. The model estimates the modulation type possible assuming DOCSIS 3.1 technology.

Return RF System Performance		Sub-Split	Mid-Split	High-Split 238
Upper Frequency	MHz	42	85	238
Homes Passed		250	250	250
HSD Take Rate		50%	50%	50%
HSD Customers		125	125	125
Desired Carrier BW	MHz	6.4	6.4	6.4
Modulation Type		2048-QAM	1024-QAM	512-QAM
Bits/Symbol		11	10	9
Number Carriers in Bonding Group		3.5	10.25	33
Max Power per Carrier Allowed in Home	dBmV	59.6	54.9	49.8
Worst Case Path Loss	dB	29.1	30.1	33.5
Maximum Return Amplifier Input	dBmV	30	25	16
Actual Return Amplifier Input	dBmV	15	15	15
Assumed Noise Figure of Amplifier	dB	7	7	7
Return Amplifier C/N (Single Station)	dB	65	65	65
Number of Amplifiers in Service Group		15	15	15
Return Amplifier C/N (Funneled)	dB	53.4	53.4	53.4
Optical Return Path Technology		uDFB	uDFB	uDFB
Assumed Optical C/N	dB	44	41	37
System C/N	dB	43.5	40.8	36.9
Desired C/N	dB	41	38	35

Table 4: DOCSIS 3.1 Capacity Prediction with Several Upstream Splits and AM Optical Technology

Though the model suggests the use of very high order modulation, it is important to know this does not account for noise conditions related to external interference or noise events. The model suggests the CNR of the channel will support very high order modulation, however these modulation formats will need to be supported from the cable modem to the burst receiver in the headend. It is to early to tell if the upstream will support as high a modulation order as the model suggests, as these systems are not available at this time.

Background on the ARRIS Upstream Model

It is important to understand what the model is and what it is not. In summary the model considers the component, which comprise the access layer for the upstream. This includes, the optical technology, distance between headend and node, coaxial electronics, passives, coaxial cable, distance, modem power, and many other factors. The diagram in figure 2, illustrates many of the measured parameters in the ARRIS Model.

The model:

- Measures the performance of the Optical and Coaxial segment
- Has flexibility to account for different network architectures and components
- Accounts for distance variations in the optical and coaxial segments
- Accounts for various service group sizes to adjust for noise funneling effects
- Accounts for noise contribution of the HFC Network
- Accounts for Attenuation
- Accounts for temperature variation in many areas
- Estimates DOCSIS 3.1 Capacity
- Model defines Operator Margin of 10 dB above the slicer for a "coded" LDPC and BCH modulation format as defined in the Figure 7 called "DOCSIS 3.0 versus DOCSIS 3.1 Modulation C/N" at the end of this section.

The model does not account for:

- Noise conditions related to external interference or noise events
- Faulty components
- Variables in Combining

QUESTION #1 ANSWER: ARE CABLE NETWORKS "LIMITED BY" THE RF VIDEO AND DATA TECHNOLOGIES? YES…

The downstream and upstream capacity is limited by the current RF video and data technologies, based on ITU-T J.83 and DOCSIS 2.0. In figure 3, the current DOCSIS based systems using J.83 Annex B shows the highest order modulation of 256 QAM provides complete coverage for all users because the network supports greater than 30 dB. The use of EuroDOCSIS J.83 Annex A in this example provides near full coverage.

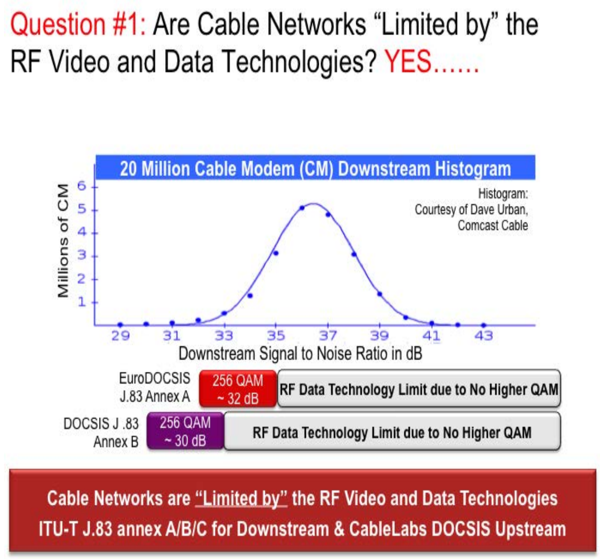

Figure 3: RF Data Downstream Technology based on J.83 is the "Limiting Factor"

Current Optical Technology and Coax Network CNR Channel Should Support DOCSIS 3.1

Return RF System Performance		Sub-Split	Mid-Split	High-Split 238
Upper Frequency	MHz	42	85	238
Homes Passed		250	250	250
HSD Take Rate		50%	50%	50%
HSD Customers		125	125	125
Desired Carrier BW	MHz	6.4	6.4	6.4
Modulation Type		2048-QAM	1024-QAM	512-QAM
Bits/Symbol		11	10	9
Number Carriers in Bonding Group		3.5	10.25	33
Max Power per Carrier Allowed in Home	dBmV	59.6	54.9	49.8
Worst Case Path Loss	dB	29.1	30.1	33.5
Maximum Return Amplifier Input	dBmV	30	25	16
Actual Return Amplifier Input	dBmV	15	15	15
Assumed Noise Figure of Amplifier	dB	7	7	7
Return Amplifier C/N (Single Station)	dB	65	65	65
Number of Amplifiers in Service Group		15	15	15
Return Amplifier C/N (Funneled)	dB	53.4	53.4	53.4
Optical Return Path Technology		uDFB	uDFB	uDFB
Assumed Optical C/N	dB	44	41	37
System C/N	dB	43.5	40.8	36.9
Desired C/N	dB	41	38	35

Cable Networks are "Limited by" the RF Data Technologies of CableLabs DOCSIS 2.0 Upstream

Figure 4: RF DOCSIS 2.0 Upstream Technology is the "Limiting Factor"

The key finding as seen in figure 3, is that DOCSIS J.83 based systems and as seen in figure 4, the DOCSIS 2.0 upstream that the limiting factor in the RF technology and not the cable plant.

QUESTION #1 SOLUTION: MODERNIZE RF DATA TECHNOLOGY WITH A NEW PHY LAYER: DOCSIS 3.1

Question #1 Solution Summary

This paper proves that the cable access network is now limited by ITU-T J.83 technology for the downstream and the DOCSIS 2.0 technology for the upstream. Theses technologies were defined as much as 15 years ago and by today's standard has low order modulation formats and an old FEC.

DOCSIS 3.1 to the rescue! The future use of DOCSIS 3.1 has four core features that will allow the MSO to maximize the network capacity or b/s/Hz. DOCSIS 3.1 key features:

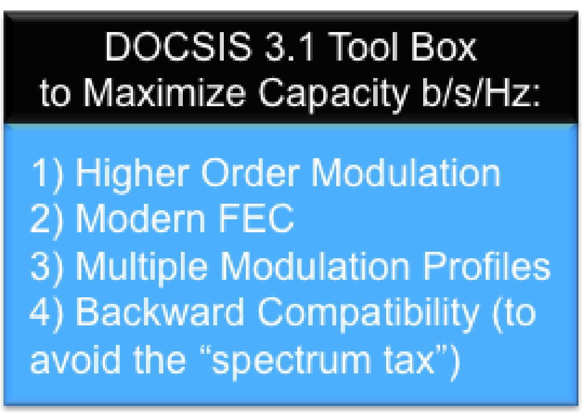

Figure 5: DOCSIS 3.1 Tool Box to Maximize Capacity b/s/Hz

As illustrated in the downstream figure 3 below and the upstream analysis as shown in the figure 4, that CNR of the channel could support higher modulation if available.

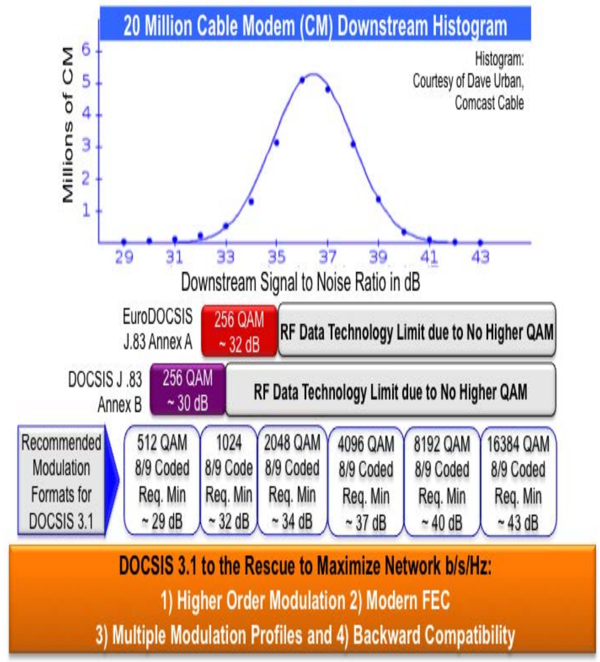

Figure 6: DOCSIS 3.1 Enables the MSOs to Maximize Spectrum Increasing b/s/Hz

The above figure 6, illustrates the point that the current RF data technology was the limiting factor preventing the MSOs to increase spectral capacity. This also suggest the use of higher order modulation to obtain more capacity is possible with DOCSIS 3.1 over the "existing" Optical and Coaxial network. However, not all users can use the same order modulations and the introduction of the use of multiple modulation profiles (MMP) is important. The use of MMP will allow groups of users the ability to reach the highest order possible, so that the network as a whole may be optimized and to maximize capacity and b/s/Hz. The use of backward compatibility allows spectrum to be shared between legacy cable modems and new modems, which support both legacy and new DOCSIS technology, avoiding a spectrum tax.

The existing cable network downstream and upstream performance can support higher order modulation formats than those available today. The support of higher order modulations with the existing network may not be ubiquitous across the MSO footprint or even within a serving group as some segments of the network will differ in performance.

Question #1 Solution Details

This section will provide details on the recommendation ARRIS made to the industry beginning in February of 2011 through February 2012, which defined the core features of later became DOCSIS 3.1. In May 2012, technology leaders from Cisco, ARRIS, Motorola, and Intel published a "Joint Supplier" paper defining with further detail the features set of what later became known as DOCSIS 3.1. These two (2) white papers, 1) ARRIS and 2) the Joint Supplier Paper are cited below and referenced in this section:

1. An ARRIS published paper in the Proceedings of the March 2012 Canadian SCTE Show titled: "Next Generation - Cable Access Network (NG-CAN), Examination of the Business Drivers and Network Approaches to Enable a Multi-Gigabit Downstream and Gigabit Upstream DOCSIS Service over Coaxial Networks" Authors: Michael J. Emmendorfer, Scott Shupe, Dr. Derald Cummings, Dr. Tom Cloonan and Dr. Frank O'keeffe.
2. The Joint Supplier (Cisco, ARRIS, Motorola, and Intel) published paper in the Proceedings of the May 2012 NCTA Cable Show titled "Mission is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS", Authors J. Chapman, M. Emmendorfer, R. Howald, & S. Shulman

The ARRIS and Joint Supplier proposals core recommendation was to modernize the DOCSIS PHY layer to increase the b/s/Hz. It proposed that DOCSIS 3.1 should:

- Use OFDM and OFDMA
- Expand the modulation orders to 4096 QAM for the downstream and the upstream.
- Add a new error correction technology or FEC for the Downstream and Upstream to include:
  - Outer FEC: Bose-Chaudhuri-Hocquenghem (BCH) codes
  - Inner FEC: Low-density parity-check (LDPC) codes
- Support for backward compatibility

The benefits of the new FEC included:

- Use of higher order modulations in similar SNR environment
- As measured against DOCSIS Upstream using A-TDMA the use of DOCSIS 3.1 with LDPC and BCH may enable a 2-order modulation rate increase in the same SNR environment, 64 QAM moves to 256 QAM.
- As measured against EuroDOCSIS Downstream using J.83 annex A the use of DOCSIS 3.1 with LDPC and BCH may enable a 2 order modulation rate increase in the same SNR environment, 256 QAM moves to 1024 QAM.
- As measured against DOCSIS Downstream using J.83 annex B the use of DOCSIS 3.1 with LDPC and BCH may enable a single order modulation rate increase in the same SNR environment, 256 QAM moves to 512 QAM.

This list of proposed features were part of the ARRIS and Joint Supplier Proposals:

- Add Downstream OFDM (Orthogonal Frequency-Division Multiplexing)
- Add Upstream OFDMA (Orthogonal Frequency-Division Multiple Access)
- Add Error Correction Technology LDPC (Low-Density Parity-Check)
- Backward Compatibility (as opposed to Coexistence) To Avoid Spectrum Tax - Leverage DOCSIS MAC across legacy SC PHY & new OFDM PHY

- Leverages every DOCSIS 3.0 MHz placed in service prior to D3.1 as part of one shared bonding group and ultimately “one network”
- Multi-gigabit Downstream and Gbps + Upstream Data Capacity
- Deployment and migration strategy to leverage existing HFC actives and use existing passives even if spectrum moves above 1 GHz and split changes
- Upstream Spectrum: Use Mid-split and/or High-split instead of Top-split
- Downstream Spectrum: 1) extend above 1 GHz (1.1 – 1.2 GHz) with existing Passives and 2) change passives when needed to support 1.7 GHz
- Extend Downstream and Upstream Modulation formats (4096 QAM)
- Continue using the Advanced DOCSIS MAC (2D Scheduler & Service Flows)
- Business Services Support (larger frame size, L2 encapsulation, etc..).
- Support for MEF bandwidth profiles in both directions and latency targets

The adoption of higher modulation formats in DOCSIS 3.1 will increase b/s/Hz. A key finding is the use of DOCSIS 3.0 Single Carrier Reed Solomon versus OFDM using LDPC may allow two (2) orders of modulation increase. In figure 7, the major takeaway from the table is the use of a stronger error correction code will allow LDPC to operate in the same carrier to noise environment as Reed Solomon but LDPC may use two orders of modulation higher.

The table uses red arrows to illustrate the corresponding Reed Solomon modulation and C/N to the OFDMA LDPC modulation format, which shares the same C/N dB. The percentage of gain is measured using the SC Reed Solomon data rate for a given modulation and the uses of two order of modulation increase allowed by using LDPC.

For example, in the table SC Reed Solomon b/s/Hz of QPSK is measured against OFDMA LDPC using 16-QAM, the percentage of gain in b/s/Hz 89%. As expected the percentage of gain will decrease as modulation increases, for example moving from 256-QAM to 1024-QAM is a smaller gain, than moving than the doubling of QPSK to 16-QAM.

The table estimates the use of OFDMA and the MAC layer bit rate in a given modulation order. This is used in the models below, which determine based on SNR for a given modulation may be selected. The model will select the highest modulation order supported if any changes are made to the inputs.

Modulation and Error Correction Comparison							MSO C/N Target
Modulation	SC Reed-Solomon MAC Layer Capacity Per MHz	Reed Solomon C/N Target (dB)	OFDMA MAC Layer Capacity Per MHz	LDPC C/N Target (dB)	Percentage of b/s/Hz Improvement of LDPC over RS		DOCSIS NG LDPC Operator Desired C/N Target (dB)
QPSK	1.229	10	1.280	4	N/A		14
8-QAM	1.843	13	1.921	7	N/A		17
16-QAM	2.457	16	2.561	10	108%		20
32-QAM	3.071	19	3.201	13	74%		23
64-QAM	3.686	22	3.841	16	56%		26
128-QAM	4.300	25	4.481	19	46%		29
256-QAM	4.914	28	5.121	22	39%		32
512-QAM	5.528	31	5.762	25	34%		35
1024-QAM	6.143	34	6.402	28	30%		38
2048-QAM	6.757	37	7.042	31	27%		41
4096-QAM	7.371	40	7.682	34	25%		44
All Mbps/MHz with the PHY Layer and MAC Layer Overhead Removed							MSO Adjustable
Single Carrier Reed-Solomon MAC Layer Capacity with 86 % Coded							
OFDMA use of more Error Correction in LDPC reduces the b/s/Hz when comparing similar modulation format of Reed Solomon							
OFDMA calculations use LDPC with 5/6 coded to achieve a 6 dB Target to Operate 2 Orders of Modulation Increase over RS							
DOCSIS NG LDPC Operator Desired C/N Target is set at 10 dB above LDPC and aimed to suggest a value that if met a desired modulation may be used							
All values are estimates and may vary based on vendor implementation and operators networks, some conditions may require different C/N targets							
All Values assume BER of 10^-8							
Percentage of b/s/Hz Improvement of LDPC over RS (Assuming 2 Order Modulation Increase)							

Figure 7 – DOCSIS 3.0 versus DOCSIS 3.1 Modulation C/N and Capacity Estimates

QUESTION #2
ARE CABLE NETWORKS "LIMITED BY" THE FTTN OPTICAL TECHNOLOGY?

The optical layer will be examined in this section. The paper will only examine the return path optical technologies and performance attributes. The optical transport return path technologies include: Amplitude Modulation (AM), commonly referred to as analog optics and Broadband Digital Return (BDR), which may be referred to as simply Digital Return.

This section will examine if the future capabilities of the cable access network will be limited by the fiber to the node (FTTN) optical technology. This section will examine the network capacity if we replaced the AM optics with digital optics, like those used for Broadband Digital Return?

In the section above it was proved that the RF date technology defined in the late 90's using Recommendation ITU-T J.83 Annex A/B/C, which is the basis for Cable's Digital Video and Data (DOCSIS) technologies of today is the limiting factor in maximizing b/s/Hz. A 15-year run! We now realized the RF technology limitation, which was a driver to modernize DOCSIS with a new PHY layer.

The section above examined the downstream and upstream performance and showed that more capacity could be achieved with DOCSIS 3.1 using the existing network. This proved that Amplitude Modulation (AM) optics used in today's HFC could support higher order modulation as those defined in DOCSIS 3.1. However, depending on upstream spectrum split, optical span and optic type the use of the highest order modulations defined was not possible. There could be many causes, the cable distribution network side, the size of the service group, the spectrum used, and it could be the optical technology.

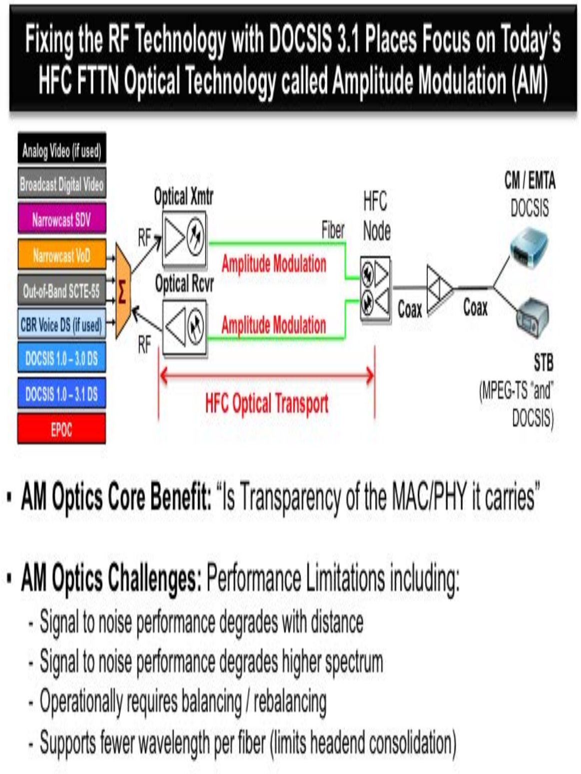

Figure 8: Overview of the Amplitude Modulation Optics

Overview of the "Current" FTTN Optical Technology

Amplitude Modulation optics is now mostly done with a Distributed Feedback (DFB) laser located in the node housing and an analog receiver located in the headend or hub. Analog return path transport is considered as a viable option for mid-split and high-split returns. Supporting short to moderate return path distances of 0-50 km with full spectrum high-split is achievable. If the wavelength is changed to 1550 nm with an EDFA, then greater distances are possible.

The analog optical return path transport presently supports up to 200 MHz loading; but typically only 5-42 MHz or 5-65 MHz is carried, depending on the distribution diplex filter split. The major benefit with analog optical return is its simplicity and flexibility, when compared with HFC style digital optical transmission. Distance is the chief challenge of analog optical transport and we will examine if support for very high order modulation, like that defined in DOCSIS 3.1 could be a factor.

Pros

The chief advantage of analog return is its cost effectiveness and flexibility. If analog return optics are in use in the field today, there is a good chance that they will perform adequately at 85 MHz; and even 200 MHz loading may be possible, if required in the future. This would allow an operator to fully amortize the investment made in this technology over the decade.

Important:

AM optics may support very high order modulation (4K & 16K QAM) though there are some restrictions mainly due to:

- Depends on the type of optic in the forward and return
- Distance, spectral loading, spectral placement in the low frequency band to achieve for the highest modulation order, and service group size (upstream)
- AM optics short distance or O-band optics will yield best performance
- Manufacture consultation is suggested to confirm performance thresholds

Cons

There are drawbacks to using analog optics. Analog DFB's have demanding setup procedures. RF levels at the optical receiver are dependent on optical modulation index and the received optical power level. This means that each link must be set up carefully to produce the desired RF output at the receiver (when the expected RF level is present at the input of the transmitter). Any change in the optical link budget will have a dramatic impact on the output RF level at the receiver, unless receivers with link gain control are used.

Also, as with any analog technology, the performance of the link is distance dependent. The longer the link, the lower the input to the receiver, which delivers a lower C/N performance.

Here is a list of challenges that Amplitude Modulation face:

- Limitations Long Distances
- Fiber distortions in AM optics can be much more disruptive to signal integrity than the coax distortions
- Many Noise Contributions in Fiber Transport Negatively Impact AM Optics
- Fiber Signal Distortions (Linear & Non-Linear)
  - Interc-hannel Crosstalk
  - Intra-channel Crosstalk
  - Non-uniform Attenuation vs wavelength
  - Non-uniform Insertion Losses vs wavelength
  - Chromatic Dispersion
  - Polarization Mode Dispersion
  - Cross-Phase Modulation
  - SBS
  - SRS
  - 4 Wave Mixing
  - Cross-Modulation
- Transmitter Electronics/Amplifier Signal Distortions (Linear & Non-Linear)
- Laser Signal Distortions (Linear & Non-Linear)
  - RIN (Relative Intensity Noise)
  - Laser Phase Noise
- Optical Amplifier Distortions (Linear & Non-Linear)
  - Spontaneous Emission Noise
  - Noise Beat components
- Photo-detector Signal Distortions (Linear & Non-Linear)

- Quantum Shot Noise
- Dark Current Noise
- Background Light Sources
- Receiver Electronics/Amplifier Signal Distortions (Linear & Non-Linear)
  - Johnson-Nyquist Thermal Noise

QUESTION #2 ANSWER:
ARE CABLE NETWORKS “LIMITED BY” THE FTTN OPTICAL TECHNOLOGY?
“NOT NOW BUT IN FUTURE YES”

We have modeled the network architecture using DOCSIS 3.1 and keeping all other coaxial conditions the same, while only changing the Amplitude Modulation (AM) to account for distance variation. The table illustrated the AM and BDR optical constrains that was change, this single change will impact the performance of the system. The change in AM optical performance will be the “limiting factor” in limiting the highest order of modulation of the future DOCSIS 3.1 system. The paper will prove the it is not just the DOCSIS 3.1 advance FEC of LDPC and Higher Order Modulation that allows for gain and/or C/N gain it is also the Optics and especially for High-Split. The table below, figure 9, illustrates the difference in performance of the return path amplitude modulation DFB optical technology versus the Broadband Digital Return (BDR). In this paper we have select one AM optical technology, which is listed in the model with the performance to support 40 km to 50 km. We could have selected a short span from the fiber node that would have yielded better results. We chose this distance, as this would likely cover 80% of all MSOs HHP.

Optical Segment Characterization			Sub-Split	Mid-Split	High-Split 238
Upper Frequency		MHz	42	85	238
AM Optics Uncooled DFB (0-25 km) single wavelength	uDFB	dB	48	45	41
AM Optics Uncooled DFB (25-40 km) single wavelength	uDFB	dB	44	41	37
AM Optics Cooled DFB DWDM (0-50 km) single wavelength	cDFB	dB	48	45	41
AM Optics Uncooled DFB (0-25 km) multi-wavelength 8	uDFB	dB	44	41	37
AM Optics Uncooled DFB (25-40 km) multi-wavelength 8	uDFB	dB	40	37	33
AM Optics Cooled DFB DWDM (0-50 km) multi-wavelength 16	cDFB	dB	44	41	37
BDR (Not Impacted by Length "or" Multi-wavelength)	BDR	dB	48	48	48

Figure 9: Optical Technology Choices

The model assumed the use of an Amplitude Modulation (AM) Optic using Uncooled DFB for use up to 40 km assuming a single wavelength. Then the model assumed the use of DOCSIS 3.1 with all of the new PHY layer improvements, such as OFDMA, a pair of errors correction technology LDPC inner code, BCH outer code, and the expansion in the available Modulation order up 4096 QAM. The model estimated Sub-split DOCSIS 3.1 modulation at 2048 QAM, Mid-split at 1024 QAM, and High-split at 512 QAM, as seen in the figure below. The model shows different modulation support depending on split option.

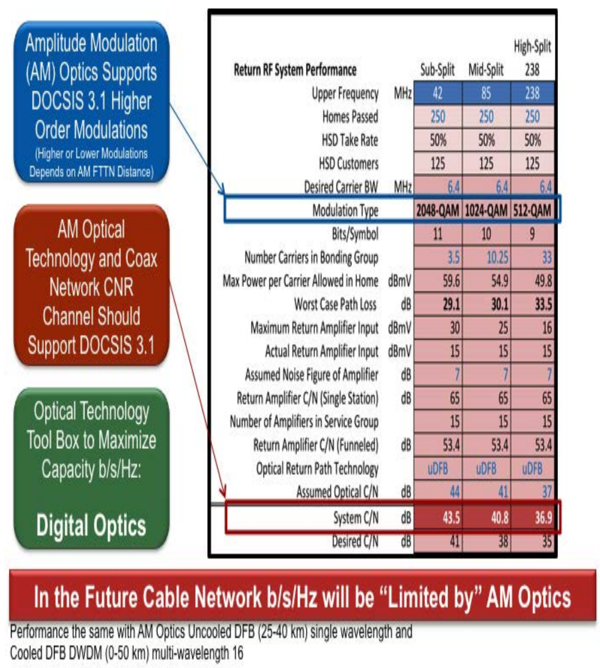

Return RF System Performance		Sub-Split	Mid-Split	High-Split 238
Upper Frequency	MHz	42	85	238
Homes Passed		250	250	250
HSD Take Rate		50%	50%	50%
HSD Customers		125	125	125
Desired Carrier BW	MHz	6.4	6.4	6.4
Modulation Type		2048-QAM	1024-QAM	512-QAM
Bits/Symbol		11	10	9
Number Carriers in Bonding Group		3.5	10.25	33
Max Power per Carrier Allowed in Home	dBmV	59.6	54.9	49.8
Worst Case Path Loss	dB	29.1	30.1	33.5
Maximum Return Amplifier Input	dBmV	30	25	16
Actual Return Amplifier Input	dBmV	15	15	15
Assumed Noise Figure of Amplifier	dB	7	7	7
Return Amplifier C/N (Single Station)	dB	65	65	65
Number of Amplifiers in Service Group		15	15	15
Return Amplifier C/N (Funneled)	dB	53.4	53.4	53.4
Optical Return Path Technology		uDFB	uDFB	uDFB
Assumed Optical C/N	dB	44	41	37
System C/N	dB	43.5	40.8	36.9
Desired C/N	dB	41	38	35

Figure 10: Support of High Order Modulation Varies with Spectrum Split

The model determines the modulation format based on the System C/N as shown in the figure above and highlight in red and with a red box. The model will select the highest modulation order supported based on the System C/N.

To verify if the Amplitude Modulation optical technology is the limiting factor this parameter will be changed in the model as seen in table 5. This single parameter was changed swapping the Amplitude Modulation (AM) optical technology with a Digital Optic using Broadband Digital Return (BDR). This single change may account for two (2) to three (3) orders of the modulations increase over use of AM optics when considering the high-split spectrum band. The sub-split and mid-split will not see as much of a gain because of the spectrum location and load is much smaller bandwidth block.

Return RF System Performance		Sub-Split	Mid-Split	High-Split 238
Upper Frequency	MHz	42	85	238
Homes Passed		250	250	250
HSD Take Rate		50%	50%	50%
HSD Customers		125	125	125
Desired Carrier BW	MHz	6.4	6.4	6.4
Modulation Type		4096-QAM	4096-QAM	4096-QAM
Bits/Symbol		12	12	12
Number Carriers in Bonding Group		3.5	10.25	33
Max Power per Carrier Allowed in Home	dBmV	59.6	54.9	49.8
Worst Case Path Loss	dB	29.1	30.1	33.5
Maximum Return Amplifier Input	dBmV	30	25	16
Actual Return Amplifier Input	dBmV	15	15	15
Assumed Noise Figure of Amplifier	dB	7	7	7
Return Amplifier C/N (Single Station)	dB	65	65	65
Number of Amplifiers in Service Group		15	15	15
Return Amplifier C/N (Funneled)	dB	53.4	53.4	53.4
Optical Return Path Technology		BDR	BDR	BDR
Assumed Optical C/N	dB	48	48	48
System C/N	dB	46.9	46.9	46.9
Desired C/N	dB	44	44	44

Table 5: Swap AM Optics for BDR and Measure the Results

The use of BDR optics provides more operating margin and higher b/s/Hz because the assumed performance of BDR is better than AM optics. In the case of sub-split and mid-split covering shorter distances or with a cooled DFB optics performance may be an near parity with BDR. The move to high-split spectrum is when in all cases the use of BDR is better than AM optics.

QUESTION #2 SOLUTION: MODERNIZE OPTICAL TECHNOLOGY: DIGITAL OPTICS

Question #2 Solution Summary

In the future will the capability of the cable access network to increase b/s/Hz be "limited by" the fiber to the node (FTTN) optical technology? Yes, however the performance of AM optics when used for sub-split and mid-split may perform at near parity against digital optics depending greatly on both distance and AM optical selection.

In the table 5 above and figure 11 below, the use of AM optics will enable higher order modulation to support DOCSIS 3.1. However to maximize DOCSIS 3.1 and remove the optical layer from becoming the limiting factor the move to digital optics in some cases will allow full support of the highest order modulations. In figure 11, is a side-by-side comparison of these findings.

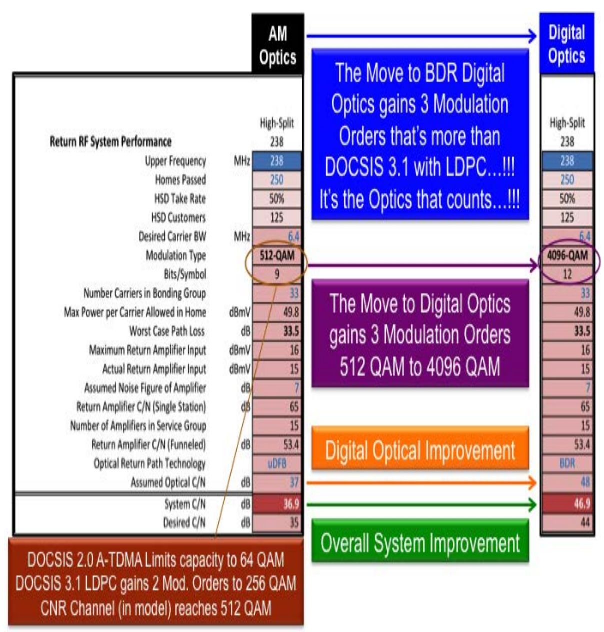

Figure 11: High-split AM Optics versus Digital Optics

The Paper Found New Key Findings with the Use of Digital Return

1. Digital Optics Maximizes Overall System Performance in terms of b/s/Hz by enabling 2 to 3 more modulation order over AM optics when considering high-split (sub-slit and mid-split the gain is smaller)
2. To maximize DOCSIS 3.1 the optical link will need to be digital for high-split
3. The use of BDR style digital optics places of only the lowest layer of the PHY in the node known as the ADC (analog-to-digital converter).
4. This places the absolute least amount of the PHY in the node to enable use of digital optics, minimizing functionality in the outside plant.

As stated above and shown in figure 11, this paper proves that there are new drivers for use of Broadband Digital Return to maximize overall system performance.

Broadband Digital Return is better than AM Optics because:

1. Digital Optics has better Performance in the Optical Segment (when compared to AM optics)
2. Signal to noise performance does not degrade with distance
3. Signal to noise performance does not degrade with return path increase in spectrum and loading
4. Better BER performance in the presence of Fiber-induced noise…(due to ability to correct bit errors)

It's the Optics!!! HFC Digital Return Matters

Question #2 Solution Details about Digital Return Path

This section examines the overall use of digital optical technology as well as the details of broadband digital return (BDR). The digital return approach is "unaware" of the traffic that may be flowing over the spectrum band of interest. It simply samples the entire band and performs an analog to digital conversion continuously, even if no traffic is present. The sampled bits are delivered over a serial digital link to a receiver in the headend or hub, where digital to analog conversion is performed and the sampled analog spectrum is recreated.

Pros

There are a number of advantages to the digital return approach. The output of the receiver is no longer dependent on optical input power, which allows the operator to make modifications to the optical multiplexing and de-multiplexing without fear of altering RF levels. The link performance is distance independent – same MER (Modulation Error Ratio) for 0 km as for 100 km, and even beyond. The number of wavelengths used is not a factor since on/off keyed digital modulation only requires ~20dB of SNR; thus fiber cross-talk effects do not play a role in limiting performance in access-length links (<160 km)

The RF performance of a digital return link is determined by the quality of the digital sampling, rather than the optical input to the receiver; so consistent link performance is obtained regardless of optical budget. The total optical budget capability is dramatically improved since the optical transport is digital. This type of transport is totally agnostic to the type of traffic that flows over it.

Summary of Digital Optics Drivers:

- **Digital Optics has better Performance in the Optical Segment (when compared to AM optics)**
    - Signal to noise performance does not degrade with distance
    - Signal to noise performance does not degrade with return path increase in spectrum and loading
    - Better BER performance in the presence of Fiber-induced noise…(due to ability to correct bit errors)
- **Digital Optics Maximizes b/s/Hz of the Coax Segment**
    - Allows for Higher Order Modulations to be used in the Coax Segment
    - Digital Optics enables use of Higher Order Modulation at any spectrum band (limited only by the RF segment)
    - Ability to make tradeoffs between Bandwidth & SNR
- **Digital Optics Improves MSO Network Operations**
    - "Set it and forget it" – technician and maintenance friendly
    - "Set it and forget it" – (vs. continual leveling & adjustments of Analog Optics for QAM Overlay Systems)
    - Digital Optics simplifies installation
    - Digital Optics can provide optical link monitoring in both directions
    - Digital Optics can provide protection switching protocols for use in case of fiber cuts (with alarms)
    - Supports redundancy over uneven lengths/longer lengths
    - Removing RF from the head-end… less complex cabling
    - Removing RF from the head-end… may save powering costs
    - Less rack-unit space in the head-end per Mbps
    - Pairs well with "fiber deep" architectures, enables "service group aggregation"
- **Digital Optics Maximizes the Optical Segment**
    - More lambdas can be packed more closely together on the fiber… important where fiber count is insufficient due to lots of node splits
    - Longer fiber reach (this can help MSOs as they plan for more head-end consolidation as well e.g. amplified or repeated)
    - Some Digital Optics solutions allow for changes in the spectrum band (optical link and spectrum band are unbounded)
- **Optical Costs**
    - Pluggable optics for less costly inventory
    - Ability to ride the Ethernet optics pricing curve
    - Likely Lower cost per Mbps… due to improved spectral efficiency & Ethernet

Cons

The chief drawback to digital return is the fact that nearly all equipment produced to date is designed to work up to 42 MHz. Analog receivers are not useable with digital return transmissions. Further, the analog-to-digital converters and digital return receivers aren't easily converted to new passbands. It requires "forklift upgrades" (remove and replace) of this optics when moving to 85 MHz and 200 MHz return frequencies. There is currently no standardization on the digital return modulation and demodulation schemes, or even transport clock rates.

Another chief drawback to digital return is the Nyquist sampling theorem. It requires a minimum sampling rate, $f_s > 2B$ for a uniformly sampled signal of bandwidth, B Hz. For n-bit resolution, this requires a Transport Clock frequency >2nB. It is assumed that the higher the transport clock, the more costly it is. And with higher clock speed, there is more fiber dispersion, which sets an upper limit on transport rate! This causes some practical limitations as to how high the return spectrum can cost effectively reach when considering digital return.

The key points about Nyquist Sampling are captured below.

Nyquist Sampling Theorem governs the minimum sampling rate

- Minimum sampling frequency must be at least twice the frequency width of the signal to be digitized (at least and most suggest a 2.5 times the sampling rate)

Nyquist Theorem causes some practical limitations

- Higher speed A/D converters typically have less Effective Number of Bits (ENOB), translating to decreasing performance at increasing clock speeds for a fixed number of bits

QUESTION #3
MOVING FROM AM OPTICS TO DIGITAL OPTICS FOR FTTN WILL FORCE US TO PLACE THE PHY OR MAC/PHY ACCESS LAYER FUNCTIONS IN THE NODE

As stated in the section above there is benefits for using digital optics and there may also be some drawbacks, such as placing more functionality in the outside plant.

Moving from AM Optics to Digital Optics for FTTN will force us to place PHY or MAC/PHY Access Layer Functions in the Node. What stays in the headend and what moves to the node? The industry will need to define a new access network architecture supporting digital connections between headend and fiber node. This new access network architecture will redefine the CCAP architecture and other headend platforms (e.g. Digital Optical Platforms) as well as the node platforms.

<u>Overview of Current and Future FTTN Optical Technology</u>

The optical layer and the relationship to the remote access layer architecture will be examined in this section.

Today, the two technologies used in optical transport for the return include Amplitude Modulation (AM) and Broadband Digital Return (BDR) as review in the preceding section. The Broadband Digital term and current application is tied to the return path, however this may be used for the forward path as well.

Broadband Digital places the lowest layer of the physical (PHY) layer called the PMD (Physical Medium Dependent) in the Node. The PMD layer of the PHY is where the ADC/DAC (Analog-to-Digital or Digital-to-Analog) functions take place.

The FTTN technology and architecture for HFC has always retained one core function --- transparency of the underlining MAC/PHY technologies that travels through it. The transparency of the RF MAC/PHY technologies was possible because of the optical FTTN technology used to include either Amplitude Modulation optical technology or Broadband Digital.

In the future we need to consider the possibility of moving the IP/Ethernet transport past the HE/Hub locations to the node. We will examine what we are referring to as a new class of cable FTTN architecture called Digital Fiber Coax (DFC). The use of DFC may augment the existing HFC media conversion class of architecture that has been deployed for about two decades. We are suggesting that there are really two different Fiber to the Node (FTTN) architecture classes for Cable Networks. These will utilize FTTN and coax as the last mile media, but this is where the similarities will stop.

To simply summarize, the Two Different Cable FTTN network architecture classes are:

- HFC is a “Media Conversion Architecture”
- DFC is a “PHY or MAC/PHY Processing Architecture”

These new FTTN technologies and architectures have or will emerge, that if implemented “may” remove this transparency.

Should the cable industry change the definition of HFC to mean multiple functions “or” define a new term(s) for this fundamentally different Class of FTTN Network Architecture that use Digital Optics to/from the node and places PHY or MAC/PHY functions in the node?

“Two (2) Different” Fiber to the Node (FTTN) Architecture Classes for Cable Network Emerging

In this section, we describe the functions of several approaches within class of fiber to the node. The following figures will aid in aligning the definitions with the list of functions, please refer to figures 12 through 17, with emphasis on figure 17.

Hybrid Fiber Coax (HFC) Class of FTTN

1. Amplitude Modulation uses Media Conversion (Optical-to-Electrical or Electrical-to-Optical) allowing for transparency of the RF MAC/PHY technologies. This is what we have using for decades.

Digital Fiber Coax (DFC) Class of FTTN

1. **Remote PMD (R-PMD)** (Physical Medium Dependent) Remotes the ADC/DAC (Analog-to-Digital or Digital-to-Analog) Allowing for transparency of the RF MAC/PHY technologies. This is used in some networks today in the return, called Broadband Digital Return (BDR). The use of BDR places the ADC in the node and there is a corresponding DAC in the headend. This type of architecture could be called Remote Physical Medium Dependent (R-PMD), because this layer of the PHY is where the DAC/ADC reside. These are proprietary solutions today but could easily be standardized. We are suggesting the term Remote PMD because this better defines the remote PHY layer that is placed in the node. This may enable a standards based return and forward path technologies using ADC and DACs in the node. This approach is the “only” Remote PHY architecture that maintains the transparency of the underlining MAC/PHY technologies that travels through it and uses digital optics.

2. **Remote Lower PHY** (RL-PHY) Remote Lower PHY is placed in the node where constellation symbols or groups of constellation symbols are received for modulation. This represents the Modulation functions and is sometimes called Remote Mod.

3. **Remote PHY (R-PHY)** This places the full PHY layer including the FEC, symbol generation, modulation, and DAC/ADC processing in the node.

4. **Remote Lower MAC (RL-MAC)** The Lower MAC functions for scheduling and the entire PHY functions are placed in the node.

5. **Remote CCAP (R-CCAP)** Places the entire upper and lower MAC and PHY layer functions in the node. This places the CMTS, Edge QAM and CCAP functions into the node.

Illustration of the Current and Future FTTN Optical Technology

The figures in the sections represent the high-level functions and technology placement in the headend and node.

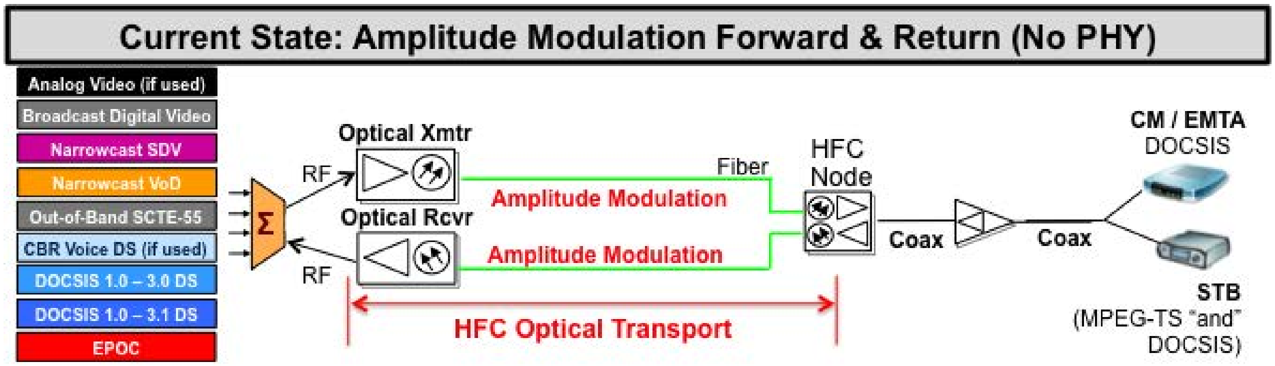

Figure 12: HFC Amplitude Modulation Forward and Return

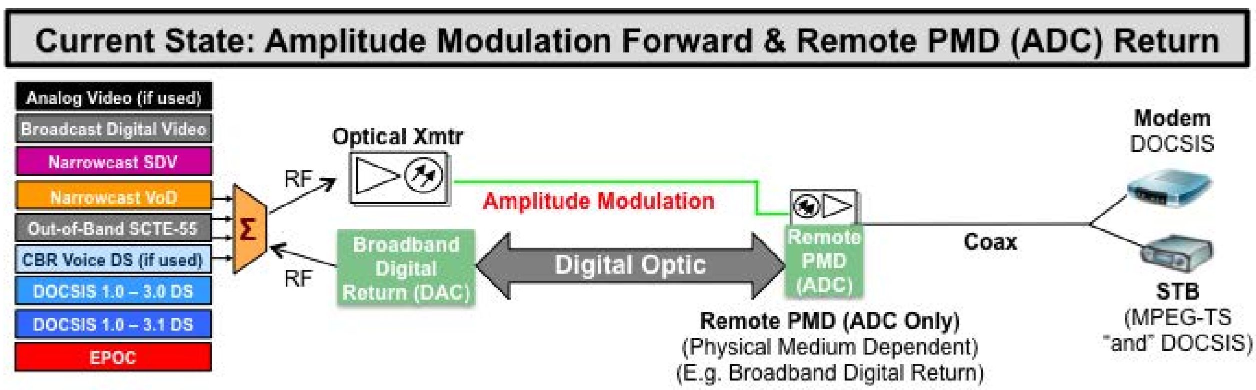

Figure 13: HFC Amplitude Modulation Forward and DFC Broadband Digital Return (BDR) (a new term could be Remote PMD)

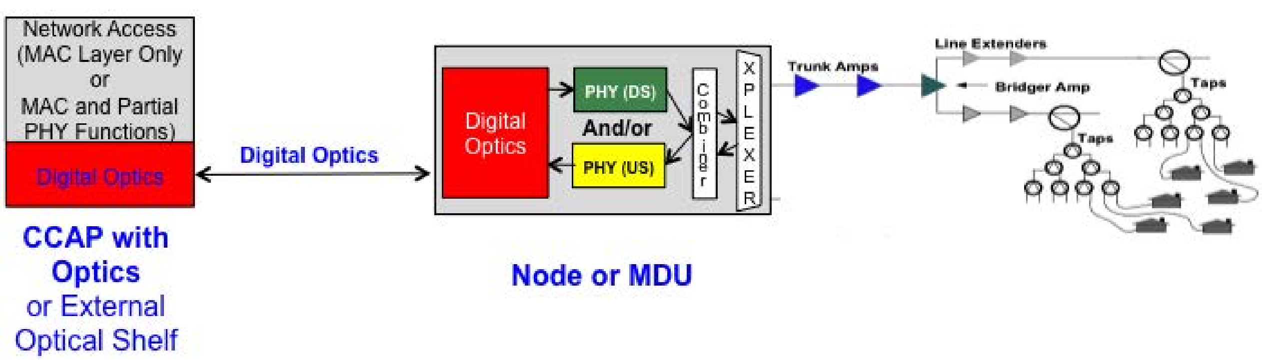

Figure 14: Digital Fiber Coax – Remote PHY Layer Options

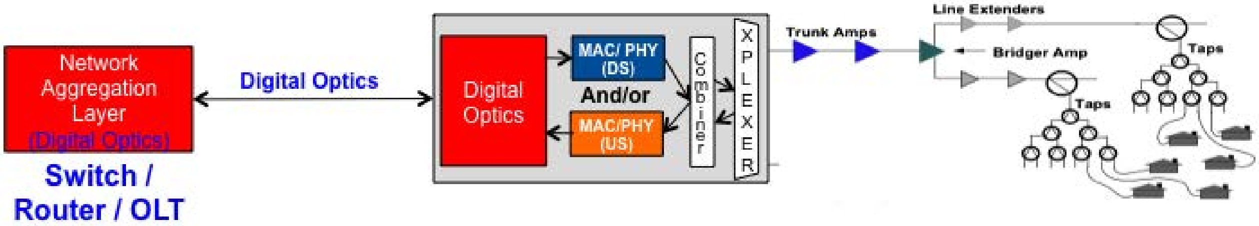

Figure 15: Digital Fiber Coax – Remote PHY Layer Options

The figured below is likely a first of its kind. This is meant to align cable technologies to the OSI reference model. The technologies examined include DOCSIS 3.0 and Edge QAM functions to the left which both use Recommendation ITU-T J.83 as the Physical Layer. The right side of the figure is an attempt to define the "possible" framework for DOCSIS 3.1 currently in development. These figures are based on the DOCSIS specifications, ITU-T J.83-B, and DVB-C2. This is aimed to help show the functions of the Remote Access Layer Architecture that may remain in the headend and that which is placed in the node.

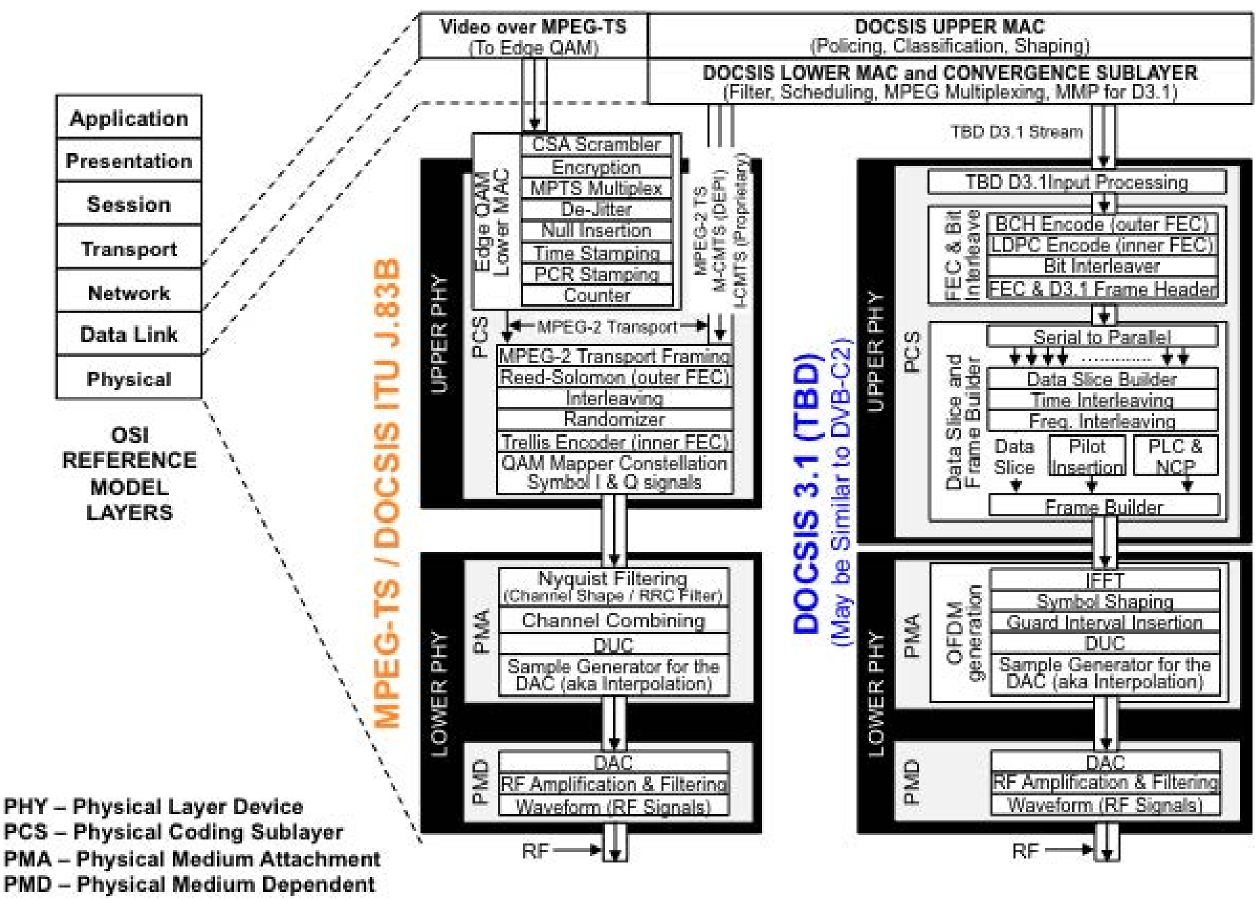

Figure 16: Digital Fiber Coax – Remote PHY Layer Options

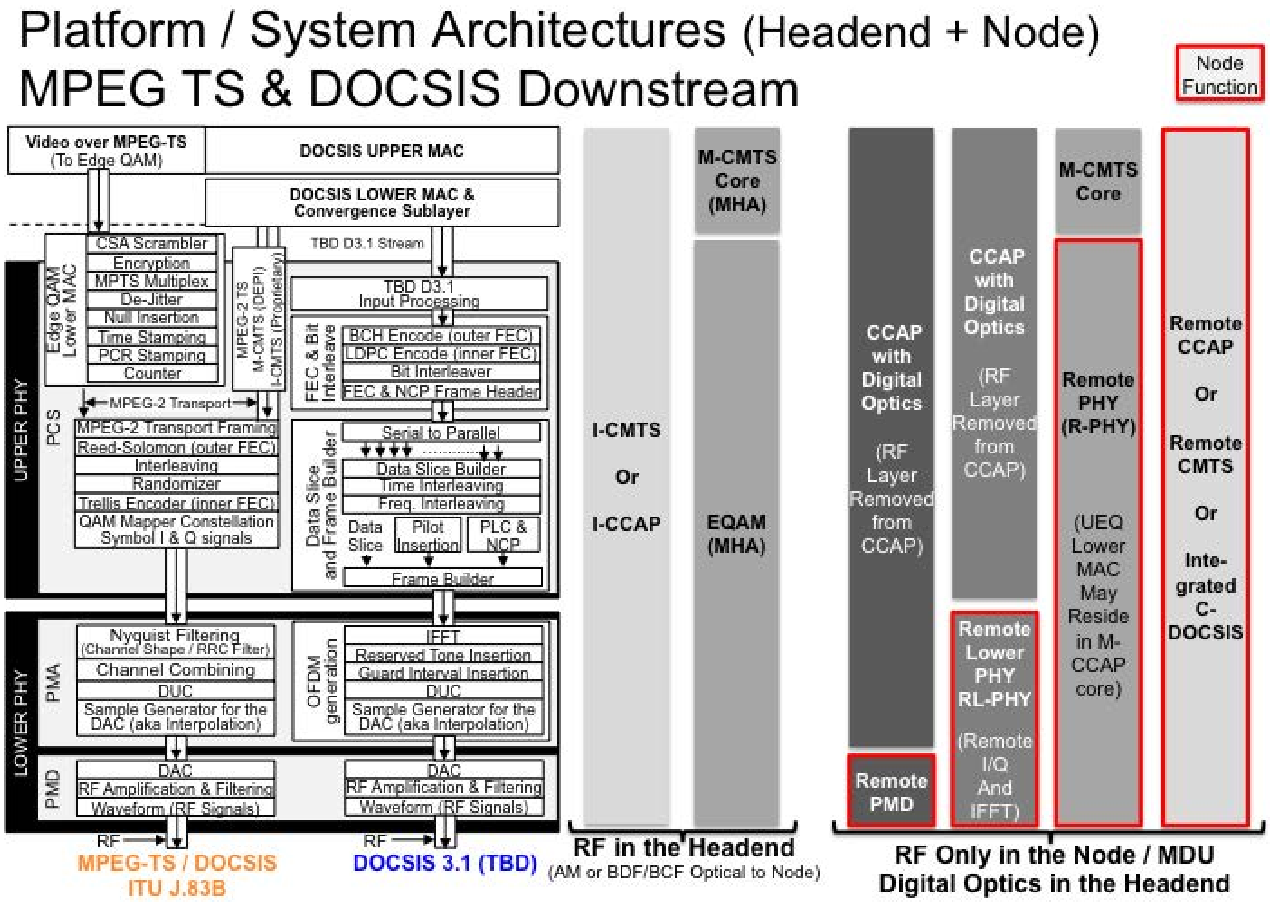

Figure 17: Digital Fiber Coax – Remote PHY Layer Options

The figure above captures the downstream DOCSIS and Edge function. The figure is intended to show the relationship with headend functions defined today and functions that will change in the headend CCAP and the node to support Remote Access Layer Architectures. The red boxes represent node functions and all align with the functions defined on the left of the figure.

CONCLUSIONS

For simplicity, we will restate the problem and provide a clear statement of the findings and conclusions.

Question #1: Are Cable Networks "Limited by" the RF Video and Data Technologies?

The capability of the cable access network to increase b/s/Hz is "limited by" the Radio Frequency (RF) Video and Data Technologies available today. These are based on Physical (PHY) Layer technologies defined by the ITU-T J.83 for the downstream used for digital video and DOCSIS downstream and the upstream based entirely on CableLabs DOCSIS 2.0 is also limiting b/s/Hz

Confirmation of Problem #1:

In the downstream, the real-world cable performance measurements of 20 million cable modems proved that the current RF Data Technologies based on DOCSIS / ITU-T J.83 is the limiting factor to maximize b/s/Hz. There is nothing the MSOs could do to increase b/s/Hz for any customer because the current J.83B technology for video and data services can operate no higher than 256 QAM.

In the upstream, the ARRIS model considered the spectrum bands of sub-split (5-42 MHz), mid-split (5-85 MHz), and high-split (5-200 MHz) and proved that the current DOCSIS upstream technology using A-TDMA 64 QAM is the limiting factor to achieving more b/s/Hz. The model showed that if a new RF Data Technology was available that defined higher modulation orders up to 4096 QAM upstream and used a better FEC that the plant SNR of the channel for Sub-split, Mid-split, and High-split could support higher order modulations, estimated at 2048 QAM, 1024 QAM, and 512 QAM respectively per split.

Solution to Problem #1:

DOCSIS 3.1 to the rescue! The use of DOCSIS 3.1 has four core features that will allow the MSO to maximize the network b/s/Hz, these features include:

1. Higher Order Modulation
2. Modern FEC
3. Multiple Modulation Profiles
4. Backward Compatibility

Question #2: Are Cable Networks "Limited by" the FTTN Optical Technology?

In the future the capability of the cable access network to increase b/s/Hz will be "limited by" the fiber to the node (FTTN) optical technology. The only technology used for forward transmission and that, which also makes up the majority of the return transmission both use an optical technology called Amplitude Modulation (AM).

Confirmation of Problem #2:

The ARRIS upstream model confirmed that the current RF Data technology using DOCSIS A-TDMA with a maximum of 64 QAM is limiting the MSOs from increasing b/s/Hz on the upstream. This model was used to confirm problem # 2. The model assumed the use of an Amplitude Modulation (AM) Optic using Uncooled DFB for use up to 40 km assuming a single wavelength. Then the model assumed the use of DOCSIS 3.1 with all of the new PHY layer improvements, such as OFDMA, a pair of errors correction technology LDPC inner code, BCH outer code, and the expansion in the available Modulation order up 4096 QAM. The model estimated Sub-split DOCSIS 3.1 modulation at 2048 QAM, Mid-split at 1024 QAM, and High-split at 512 QAM.

Then a single parameter was changed in the model, we replaced the Amplitude Modulation (AM) optical technology with a

Digital Optic using Broadband Digital Return (BDR). This single change enables the entire spectrum splits options including, Sub-split, Mid-split and High-split to be capable of reaching 4096 QAM in the upstream.

<u>Solution to Problem #2:</u>

The use of Amplitude Modulation optics can support high-order modulation for the sub-split and mid-split bands. The in the upstream enable the use of the full DOCSIS 3.1 modulation format upstream to include 4096 QAM. The high-split spectrum choice had the most gain with the use of digital optics. The downstream was not model for use of digital optics, however it is verified by the cable downstream measurements in section 1 that the downstream may support many high order modulation defined in DOCSIS 3.1

<u>Question #3: Moving from AM Optics to Digital Optics for FTTN will force us to place PHY or MAC/PHY Access Layer Functions in the Node.</u>

Moving to Digital Optical Technology for FTTN will change the access layer architecture placing PHY or MAC/PHY functions in the node. a) We need to determine what functions remain in the headend and those placed in the node. b) We need to determine the network access layer architecture if portions of the PHY or MAC/PHY are separated between the headend and the node, as this will impact the system architecture for CCAP and the connections and functions it contains in this new Digital Fiber Coax world. The paper provides an overview of the functions that may reside in the headend and node location depending on the Remote Access Architecture selection.

<u>Initial Thoughts:</u>

Placing the lease amount of functions in the node, like remote PMD or placing the ADC/DAC in the node provides the transparency of the MAC/PHY layer technologies. This is a core benefit the industry has depended on for over two decades and several service and technologies changes. It is far too early to tell which is the best approach. In the future we will publish additional research, which compare these options side-by-side.

ACKNOWLEDGEMENTS:
The authors wish to thank Brent Arnold and Frank O'Keefe of ARRIS.

REFERENCES:

[1] Recommendation ITU-T J.83 Digital multi-programme systems for television, sound and data services for cable distribution
[2] "Mission is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS", Authors J. Chapman, M. Emmendorfer, R. Howald, & S. Shulman, (Cisco, ARRIS, Motorola, and Intel) published paper in the Proceedings of the May 2012 NCTA Cable Show

Hierarchical QoS, The Next Step In Enhancing DOCSIS QoS Architecture

Alon Bernstein (alonb@cisco.com), Pawel Sowinski (psowinsk@cisco.com)
Cisco Systems, Inc.

Abstract

Robust and granular QoS support is a key factor in the tremendous success DOCSIS has achieved as a mature technology in delivery of data, voice and video services. The basic DOCSIS QoS architecture is built upon the foundation of the service flow concept that has been defined fourteen years ago as part of DOCSIS 1.1. The fundamental DOCSIS QoS framework has remained mostly unchanged since then.

Recently, as part of the industry dialog leading to CCAP and DOCSIS 3.1, cable operators have requested to apply QoS policies to aggregation of service flows in addition to individual service flow QoS. The paper presents a case for extending DOCSIS to include elements of hierarchical QoS (HQoS) technology for this purpose. While HQoS techniques have been deployed in other broadband access technologies, integration of HQoS with DOCSIS has never been attempted before. Such integration poses a number of unique technical and business challenges that deserve careful examination. Through the review of typical use case examples, the paper examines how the HQoS technology provides the business value to the cable operators in extending services to gigabit and beyond.

Finally, the paper explores a set of issues and potential solutions imperative to enable seamless integration of hierarchical QoS into CCAP deployments and to facilitate rudimentary multi-vendor interoperability, including necessary signaling protocol and elements of standard CMTS/CCAP configuration.

DISCLAIMER

The ideas described in this paper are part of Cisco's contribution to DOCSIS 3.1 specifications.

DOCSIS 3.1 specifications are still under development. The following represents authors' current thoughts on what HQoS in DOCSIS might look like, and do not represent actual decisions made regarding the final form of the specs or technology.

Anything could change.

INTRODUCTION

Brief history of DOCSIS QoS

In the late 90's cable systems were deployed without QoS. The assumption was that the 27Mbps over a 64QAM digital cable channel was more than enough bandwidth and when there is no congestion there is no need for QoS. It turns out that the above assumption was not true; TCP/IP which is the transport building block of the internet is "greedy" by nature (i.e. it attempts to fully utilize whatever pipe it has). On top of that a good percentage of cable subscribers are running greedy applications such as file sharing and servers.

QoS became a competitive issue; telecom companies came up with commercials claiming that a neighbor's activity can restrict a subscriber bandwidth. This was not a fair statement (the telecoms had their own aggregation bottleneck at the DSLAM output and in some cases those were less then 27mbps) but it did demonstrate the need to include QoS mechanisms in the cable access.

Since the need for QoS became obvious, the first version of DOCSIS (1.0) supported basic QoS. Each subscriber had a “QoS class of service” with traffic SLA (Service Level Agreement) defined per cable modem for the upstream and the downstream. While the above model helped resolve the congestion issues it clearly did not support multimedia service types for cable modems, e.g. voice/video/data.

The Requirement to support voice presented a clear and revenue-generating-reason to define QoS mechanisms that will assure voice quality even when the network becomes congested. In order to support voice and other multimedia definitions DOCSIS 1.1 changes the QoS model significantly:

1. The DOCSIS 1.0 class of service was obsoleted and instead DOCSIS 1.1 defines “unidirectional service flows” that could be separately created for the US and DS
2. New scheduling modes were defined for US to minimize latency for multimedia applications
3. The provisioned/authorized/admitted /activated states where defined for service flows

Note that the DOCSIS specifications define a behavior and not an implementation. In that spirit, this document describes how an HQoS should behave without defining how to implement it.

Current State of DOCSIS QoS

The DOCSIS 1.1 QoS model is still the base model for DOCSIS 2.0 and DOCSIS 3.0. The most significant recent update has been the aggregate QoS model for DPOE.

DOCSIS QoS Applications and Hierarchical QoSs

Hierarchical QoS allows an operator to define QoS policies on an aggregation of flows. Hierarchical QoS is defined as a strict tree structure where the physical interface capacity is typically the root (or “parent”) node. The word “strict” means that for a given child node there can be one and only one parent.

Hierarchical QoS can be implemented by means of policing, shaping and/or scheduling. Some methods include rate limiting, rate shaping and/or marking packets for Weighted Random Early Discard (WRED). Some scheduling methods include prioritization and/or Weighted Fair Queuing (WFQ).

The current DOCSIS QoS definitions do not support aggregate QoS policies, but one can argue that the aggregation of all DOCSIS service flows into a single physical channel forms a simple two-level hierarchy where the physical level (in this case a QAM channel) is the parent node and the DOCSIS service flows are the children nodes. This is shown in the figure below.

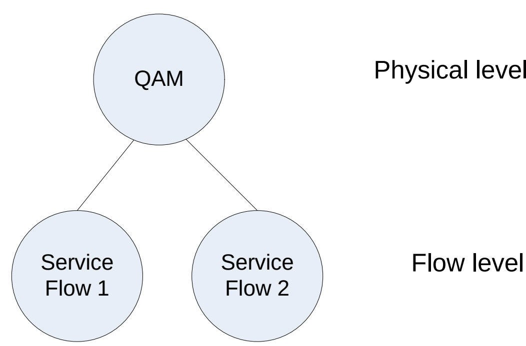

Figure 1: DOCSIS QoS model hierarchy

Hierarchy can be nested to create more complex trees. The paper will outline some of the use cases for these more complex hierarchies.

Hierarchical QoS and Fairness

Hierarchical QoS assures fairness across the children at a certain level, but not below or above. The following example can help explain the type of fairness that may be delivered in contrast to the DOCSIS 3.0 fairness model.

Fairness in the current, flat DOCSIS QoS model: if 1000 flows of equal priority share one congested interface then each flow gets 1/1000=0.1% of the bandwidth.

Fairness within Hierarchical QoS: Let's say we separate the traffic onto two virtual pipes, one for ISP A and one for ISP B, each limited to 50% of the total bandwidth, and use Hierarchical QoS:

If ISP A has only 1 active flow, and ISP B has the other 999 flows then:

- The single ISP A flow will get 1/2 = 50% the bandwidth
- Each ISP B flow will get (1/2)/999 = 0.05% of the bandwidth

In other words, Hierarchical QoS can assure fairness within the virtual pipes that it carves, but not across them.

HQoS in Edge Routers and Other Access Technologies

In the telecom world, hierarchical scheduling has been used for many years because of the multiple congestion points between the BRAS (Broadband Remote Aggregation Server) and the CPE (Customer Premise Device). Those include the interface to the DSLAM (Digital Subscriber Line Access Multiplexer) and the twisted pair to the CPE. Each one of these bottlenecks can be modeled as a "logical pipe" in a hierarchical scheduler. This use case was not needed by cable subscribers; however, as discussed in this paper, other use cases have emerged and renewed the interested in hierarchical QoS.

USE CASE REVIEW

In the initial discussions the cable operators have identified a number of use cases for HQoS. While those use cases differ in some details, such as the scaling numbers, the class of service (residential vs. business) or the types of service (data, video, and voice) in the end, the use cases boil down to two main service scenarios. These two scenarios are presented in detail below as subscriber level HQoS and service group level HQoS.

Use Case 1: Per Subscriber Aggregate QoS Controls

Recently, in order to effectively support more diverse service offerings the cable operators requested the ability to apply traffic controls, not only to individual service flows but also to groups of SFs belonging to a particular CM or IP host. The aggregate QoS limits must be enforced in addition to per Service Flow QoS treatment. The aggregate QoS enables the cable operators to offer SLAs with a simplified external structure and to more effectively compete with other ISPs which for many years have been providing internet access with hierarchically organized services.

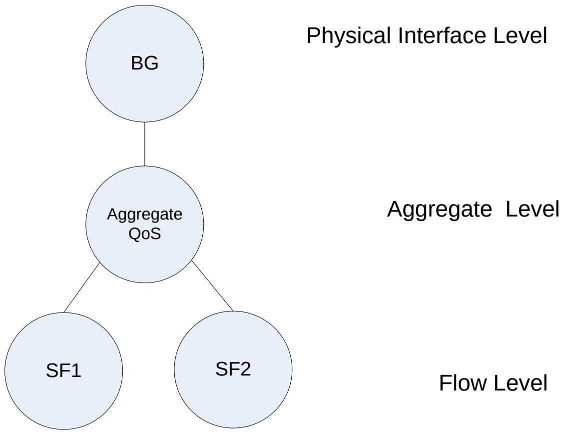

Figure 2: Use Case 1, Per Subscriber Aggregate QoS

As a typical example of use case 1, Figure **2** may be interpreted as a depiction of the QoS constructs for a downstream residential service with two QoS applications and a common, aggregate QoS policy. The individual applications are provisioned with Service Flows so that their traffic can have independent QoS treatment.

Let's consider the following settings:

- SF 1 represents a High Speed Data service with offered Maximum Sustained Traffic Rate at 20 Mb/s with Traffic Priority of 1.
- SF 2 represents a Managed Video service with offered Maximum Traffic Rate of 9 Mb/s with Traffic Priority of 5 as well as a Minimum Reserved Traffic Rate of 3 Mb/s.
- The aggregate QoS settings limit the combined traffic rate of both flows to 20 Mb/s.

The SLA which is structured with two levels of QoS controls gives the subscriber a single overall service rate at 20 Mb/s but allows the Managed Video service to operate with higher priority and with bandwidth reservation to guarantee a minimum level of QoS. When no traffic flows through the video service flow, the HSD service can use the whole 20 Mb/s. When the managed video is active, its traffic eats into the bandwidth offered to HSD service. Without aggregate QoS limit, both offered applications would run independently and potentially consume bandwidth at a higher cumulative level up to 29 Mb/s.

Use Case 2: Virtual Partitioning of a Physical Interface or a Service Group

As service offering diversifies, the cable operators would like to manage bandwidth allocated to each service from the total pool available in DOCSIS service group. Since the business model of HFC is largely relying on the concept of over-provisioning, during the periods of high usage (busy hours) the services compete for bandwidth and can negatively impact each other. This is where issues of fairness and the business model intersect. For example, a cumulative usage of bandwidth by a service with higher traffic priority can restrict bandwidth from service operating with lower DOCSIS traffic priority. It's imperative that the operators have effective tools to deal with such problems.

Today, DOCSIS 3.0 configuration offers the cable operators a convenient mechanism for this purpose. Operators may separate services by creating downstream bonding groups from distinct set of channels and appropriately "steer" the traffic belonging to each service. Such solution, based on partitioning of channel resources seems feasible considering that DOCSIS 3.0 DS channel pool in a SG scales into the tens and that each channel has relatively small capacity (38.8 Mb/s).

This method has its drawbacks as it is not applicable to upstream direction in HFC plants with low-split and mid-split. Additionally, physically partitioned service groups support lower peak traffic rates and don't share excess bandwidth. Further, with DOCSIS 3.1 such service separation scheme is no longer practical, as the channel bandwidth grows dramatically (1.7 Gb/s) and the channel count may drop down to just a few without enough granularity to effectively use them for traffic engineering.

PROPOSED DOCSIS QOS EXPANSION

Approach

When planning an expansion to an existing networking protocol, especially when it is deployed as widely as DOCSIS, one has to carefully consider a number of factors that define the overall fit of the newly added functionality into the existing architecture. These criteria, including backwards compatibility, consistency with the current methodology, multivendor interoperability as well as the ability to support incremental deployments have been contemplated when deciding the approach to introduce HQoS into DOCSIS.

In the end, we believe that the HQoS framework presented within this paper demonstrates a healthy compromise between the cable operators' requirements and how they fit into the existing DOCSIS architecture. The proposed framework includes the definition of devices' roles, key constructs, protocol signaling and common CMTS configuration.

On the other hand, the proposal does not dive into the details of implementation or internal algorithms used for queuing and scheduling. Further, the last component of the proposed framework, the CMTS configuration, may be extended in vendor-proprietary manner to meet individual vendor's needs and to provide solution differentiation.

The paper discusses HQoS primarily in the context of real-time traffic engineering and QoS policy enforcement. HQoS does impact non-real time functions such as admission control and resource management. These functions have to accommodate HQoS and the new policy controls HQoS provides. However, since these functions fall out of scope of current DOCSIS standards we feel there is no need to incorporate them in the proposed HQoS framework.

Lastly, it may be useful to note that the ideas describes in this proposal are symmetrical; they are equally applicable to traffic control for upstream and downstream directions.

The roles of CMTS and Cable Modems

HQoS is proposed as a CMTS only feature. The CMTS is responsible for all HQoS configuration and management. All aggregate QoS policy enforcement functions, including the real time traffic scheduling and queuing are performed only by the CMTS. The CMTS provides all network management capabilities necessary for status reporting related to HQoS. Cable Modems are not aware of the HQoS. CMs are required to convey HQoS information from CM configuration file into Registration Request without the need for interpretation of transported information. CMs need only implement certain QoS functions related to upstream bandwidth request policing on per SF basis only, as it is done in DOCSIS today.

New QOS Constructs

The key new construct introduced by HQoS is the Aggregate Traffic Class or the ATC. An ATC constitutes the middle point in the scheduling hierarchy; a point which is located in between service flows and physical interfaces. An ATC represents a group of service flows, or more precisely the aggregate of traffic flowing through a defined set of service flows.

As service flows, all ATCs are unidirectional; all service flows grouped into an ATC must serve the same upstream or downstream direction. In case of need for aggregate QoS policy enforcement for both directions, separate ATCs must be defined to group upstream flows and for a group of downstream flows.

Service flows may be mapped or associated with a single ATC through methods explained further in the paper. On the other hand, each ATC must be mapped onto exactly onto one physical interface to maintain hierarchical organization.

An ATC may group service flows with different QoS parameters. For example an ATC may aggregate traffic from service flows with different traffic priorities or different Maximum Peak Rates or maximum sustained rates.

ATCs have to be defined as part of the HQoS framework because they are externally visible; the information related to ATCs is exchanged in DOCSIS protocol and in the CMTS configuration.

Aggregate QoS Parameters

The CMTS enforces traffic control policy on an ATC. For this purpose, each ATC has an associated set of Aggregate QoS Parameters (AQPs). The parameters quantify the traffic control policy enforced by the CMTS. The set of proposed AQPs is listed below:

- **Aggregate Maximum Traffic Rate.** This is a mandatory parameter. It defines the maximum rate that the CMTS enforces on the aggregate of traffic flowing through the ATC. The choice of a specific algorithm for real-time enforcement of this parameter is left to vendor defined implementation.
- **Weight**. Weight controls arbitration when multiple ATCs "compete" for bandwidth of an interface. "Weight" should not be confused with "Traffic Priority". Weight is an optional parameter.
- **Aggregate Minimum Reserved Traffic Rate.** By default, an ATC will operate with the Minimum Reserved Traffic Rate which is the sum of the values of Minimum Reserved Traffic Rate parameter for each of its member SFs. This parameter is provided to allow the operator to override the Minimum Reserved Rate value "inherited" from SFs. Aggregate Minimum Reserved Traffic Rate is an optional parameter.

Note that the list of proposed AQPs does not include all QoS parameters that can be defined for individual Service Flows. Some of the excluded parameters such as Traffic Priority or Maximum Burst size can be only defined at the Service Flow level.

Two methods are proposed to provision Aggregate QoS Parameters. The first method is based on explicit inclusion of an AQP set within an ATC definition. The second, indirect method relies on creation of named AQP profiles within the CMTS's device configuration. This method resembles the existing DOCSIS mechanism for defining named Service Classes. Each AQP profile is identified by a name in the form of a string. The AQP profiles can be correlated to ATCs by name. Note, that when the AQP set is defined as part of the CMTS configuration the parameters can be augmented to include vendor proprietary extensions.

ATC Categories

This paper proposes two distinct categories of ATCs: Subscriber ATCs (SATCs) and Interface ATCs (IATCs). SATCs and IATCs differ in:

1) The purposes they fulfill
2) The selection and scaling of service flows mapped to them
3) The methods for provisioning and instantiation

Subscriber ATCs

Subscriber ATCs serve as an implementation tool for service layer agreements with two levels of QoS parameters. SATCs provide the outer QoS envelope, while service flows define QoS parameters for more granular, individual services or applications.

Each SATC aggregates traffic from a subset of unicast service flows associated with a single subscriber or a single Cable Modem. Not all service flows defined for a particular CM have to be a part of an SATC. Certain service flows may be directly mapped into physical interfaces. While Use Case 1 called for a single aggregate QoS per subscriber, the operator may be able to create more than one SATC per a cable modem. In such case, one set of CM's SFs may be mapped to one SATC while other SFs are mapped to other SATCs. As a general rule, a SF cannot be mapped to more than one ATC.

Figure **3** presents an example of a superset of Use Case 1 with five service flows and their relationship to two SATCs. In the example all five service flows belong to a single Cable Modem. SATC A includes service flows #1 and #2. SATC B consists of service flows #3 and #4. Service flow #5 is directly associated with bonding group BG Y so it is not mapped to any SATC.

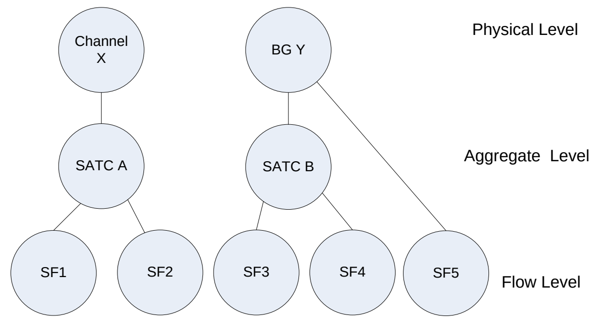

Figure 3: An example of SATCs for a single CM

How are SATCs configured? As a highly scalable edge platform, the CMTS does not maintain per subscriber configuration elements. Configuration objects that have per subscriber scaling, such as service flows, are typically configured via DOCSIS provisioning systems. Such objects are instantiated at the CMTS when a CM registers and conveys the content of its configuration file to the CMTS. SATC configuration and instantiation methods follow the practices devised for service flow provisioning. Such approach not only promotes scalability but also enables straightforward assignment of service flows to SATCs.

Operators provision SATCs by including their definition in the CM configuration file. The CM configuration file encodings are augmented to permit SATC definition. Table **1** lists the newly added CM configuration file encodings:

Attribute Name	Description
SATC Reference Number	A number identifying SATC in the CM configuration file
AQP Set or SATC AQP Profile Name	A set of scheduling parameters quantifying the QoS policy enforced by the SATC or a string which provides a reference to a named SATC AQP Profile. SATC AQP Profiles may be configured at the CMTS. The attributes: "AQP Set" and "SATC AQP Profile Name" are mutually exclusive as they provide alternative methods for provisioning of aggregate QoS policy parameters.
Service Flow Matching Method	A method by which dynamically created service flows can be matched to the SATC. The options for SF matching method are: 1. by SF Application Id 2. by SF priority range 3. by SF SCN
Service Flow Matching Criteria	A set of criteria which are dependent on the selected SF matching method.

Table 1: SATC Attributes

The service flows which are statically provisioned in the CM configuration file can be matched to SATCs by the SATC reference number as shown in

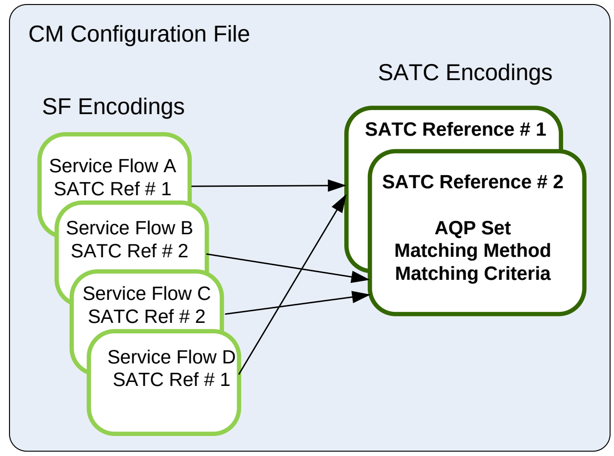

Figure **4**.

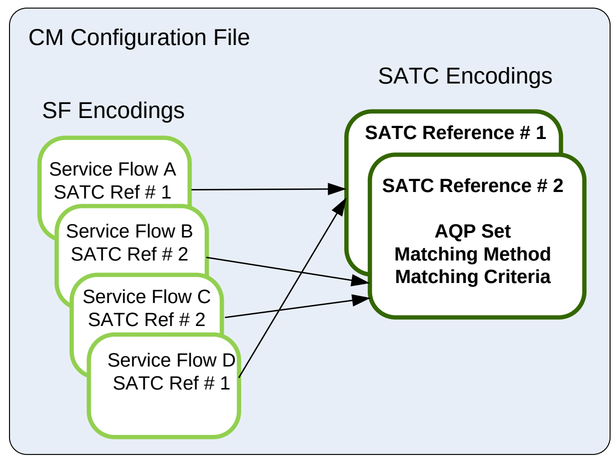

Figure 4: Correlation of SFs to SATCs in CM configuration file.

Dynamically provisioned service flows, for example those flows that are created through PCMM interface, may be matched to an SATC by one of the listed matching methods by means of the SF matching criteria.

Even though SATCs exhibit many similarities to service flows, the analogy between these constructs has limitations. Unlike DOCSIS Service Flows, SATCs don't maintain the QoS state attributes (provisioned, admitted, active), multiple QoS parameter envelopes or require associated QoS state management protocol. Cable Modems are generally not aware of SATCs.

Theoretically, the CMTS could be instrumented to create SATCs without explicit configuration; none of the identified use cases necessitates the dynamic creation of SATCs. The need for SATC QoS state management is further abated because the CMTS's admission control functions must operate at Service Flow level. Therefore the paper asserts that SATCs should be generally considered to be static objects, always present and active after a CM completes its registration.

Recently, DPoE specifications introduced into DOCSIS the concept of Aggregate Service Flows (ASFs). ASFs have been added to DOCSIS proper (MULPI) for the purpose of reserving TLV numbers. ASFs in DPoE and SATCs as proposed here for DOCSIS provide largely equivalent functionality and are provisioned in a similar way. They differ in the environment for which they have been designed and certain operational requirements. For example ASFs in DPoE may be associated with a full Service Flow QoS Parameter Set, in some cases these parameters are overridden by equivalent MEF parameters. The AQP proposed for SATC includes fewer parameters. While the authors believe that these concepts of SATC and ASF can be merged, doing so is outside of the scope of the paper.

Interface ATCs

Interface ATCs are intended to fulfill the premise of the Use Case 2. The IATCs enable the operators to virtually divide the bandwidth of service groups or physical interfaces between distinct services or users.

An example of such partitioning is shown on

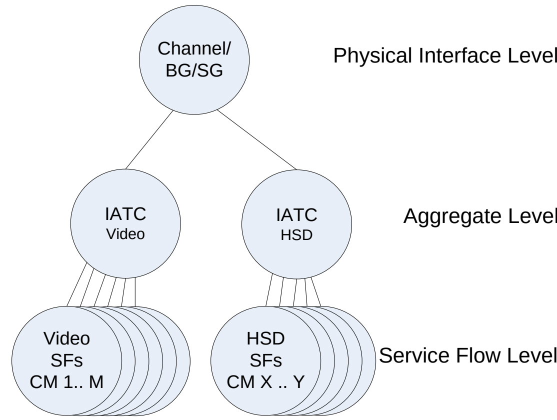

Figure **5**.

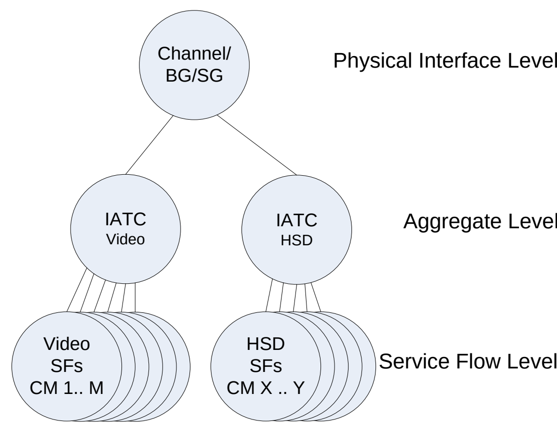

Figure 5: Interface ATCs

An IATC aggregates traffic from Service Flows belonging to multiple CMs but typically sharing some common property like application, the type of service or selected traffic priority. For example, one IATC may be created to group traffic from all service flows carrying cable operator's managed video traffic that are mapped to a particular bonding group. In the same example a second IATC may group all HSD flows mapped to the same bonding group. The aggregate QoS parameters associated with each IATC will define how the bandwidth of the bonding group is shared between managed video and HSD.

IATCs are provisioned via the CMTS configuration in a two-step process. In step one; the operator defines a number of IATC profiles. IATC profiles are identified by a name and can be used throughout the system. IATC profiles serve as templates for creation of IATC instances. Each IATC profile includes the attributes as listed in Table **2**.

Attributes	Description
IATC Profile Name	A string that uniquely identifies the IATC profile.
Aggregate QoS Set	A set of parameters defining the QoS policy enforced by the IATC.
SF Matching Method	A method by which the CMTS can match service flows (both static and dynamic) to the IATC. The following methods are proposed for SF matching: • by Application Id • by SF priority range • by SF SCN • None Note: "None" matching method may be selected when statically defined service flows in CM configuration file are explicitly matched to an IATC profile by name.
SF Matching Criteria	The set of criteria that corresponds to the configured matching method: Application Id, SCN, SF traffic priority range.

Table 2: IATC attributes

In step two of the configuration process an operator can associate any selected physical interface with one or more IATC Profiles. When more than one IATC profile is associated with an interface then the SF matching method or SF matching criteria must differ between IATC Profiles to ensure unambiguous matching decision. Not all physical interfaces must be paired to an IATC Profile. Figure **1** demonstrates an example of configuration defining the association between static physical interfaces and IATC profiles.

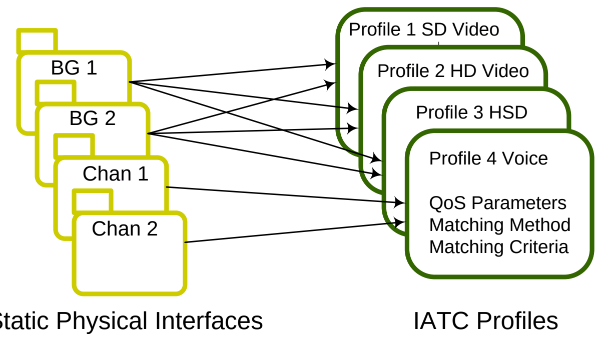

Figure 6: Static Physical Interfaces to IATC Profiles Mapping

Based on such configuration, the CMTS creates instances of IATCs for each configured pair {interface, IATC profile}. In the example shown on Figure 10, 3 IATC profiles will be instantiated for each bonding group, BG 1 and BG 2 and one IATC profile will be instantiated for each of the channels. All together the shown configuration results in creation of 8 IATCs.

The IATC provisioning method described above can be deployed for those physical interfaces that are created statically. DOCSIS allows CMTS's support for the dynamic creation of upstream or downstream bonding groups. Yet, this function is largely left to CMTS vendor definition because DOCSIS does not define a specific method or standard configuration for this purpose. This proposal takes a similar approach to the definition of HQoS over dynamically created bonding groups. We acknowledge that dynamic BGs can be associated with IATCs, but the specification of such method is left to CMTS vendor differentiation.

How are Service Flows mapped to IATCs? In the absence of H-QoS the CMTS maps Service Flows to bonding groups or individual channels. With HQoS the SF mapping process needs include one additional step: a decision whether to assign a SF to an IATC and which IATC to select. Operators will be able to control SF to IATC association via several matching methods. Those methods are defined as part of IATC configuration and listed in Table **2**.

An alternative mechanism permits association of SFs provisioned via CM configuration file to IATC profiles by name. The SF encodings in the CM configuration file are augmented with IATC name for this purpose.

IATCs typically group a subset of service flows from a service group, scaling up to hundreds of SFs. SATCs will typically aggregate traffic from a much smaller SF grouping, perhaps reaching into the teens.

The Impact on DOCSIS Protocol

The presented HQoS framework has a small impact on the current DOCSIS protocol. The scope of changes is restricted to augmentation of the CM configuration file encodings to support a few added HQoS TLVs. The CM opaquely conveys the new TLVs to the CMTS during registration. Otherwise, HQoS does not require any additional support by the CM software or hardware.

CONCLUSION

DOCSIS has been a very successful protocol. The granular QoS support in DOCSIS is part of this success, and can continue as such especially if its functions can be adapted to better serve the new range of services. This paper proposes a natural expansion to DOCSIS QoS to provide the cable operators with control over QoS policies at the aggregate level. The expansion creates tools to enable the operators to offer new and better structured services to their customers as well as to more effectively manage the allocation of one a most valuable resources at their disposal: the bandwidth of the DOCSIS part of the HFC.

ACKNOWLEDGMENTS

We would like to extend our sincere gratitude to many talented participants of the DOCSIS 3.1 MAC Focus Group and the AMP Work Group for their comments and contributions to the paper.

List of Acronyms

AQP – Aggregate QoS Parameters
ASF – Aggregated Service Flow
ATC – Aggregate Traffic Class
HQoS – Hierarchical QoS
HSD – High Speed Data
IATC – Interface ATC
MEF – Metro Ethernet Forum
QoS – Quality of Service
SCN – Service Class Name
SF – Service Flow
SG – Service Group
SATC – Subscriber ATC

References

1. DOCSIS MULPI : http://www.cablelabs.com/specifications/CM-SP-MULPIv3.0-I20-121113.pdf
2. Packet Cable MultiMedia Specification: http://www.cablelabs.com/specifications/PKT-SP-MM-I06-110629.pdf
3. AMP Technical Report: MAC Layer Design for Efficient Multi-Gigabit Transport over CATV Networks. CableLabs.

High Efficiency Video Coding (HEVC) in a Changing World—What Can MSOs Expect?

Mukta Kar, CableLabs; Yasser Syed, Comcast Cable; Munsi Haque, Consultant

Abstract

The cable operator world has been undergoing a sea-change over the last couple of years. Content is increasingly being viewed in a non-linear fashion. Service providers are not just delivering to the leased set-top box (STB), but to PCs, gaming machines, tablets, cell phones and other customer owned and managed (COAM) devices. The migration of Picture quality is not just from standard definition (SD) to high definition (HD), but also to 3D, 4K (also called UltraHD), and several resolutions in-between. STB Video Processors are changing from dedicated hardware processors to general-purpose multicore processors running video processing applications.

Amidst these changes, MPEG High Efficiency Video Codec (HEVC, also called H.265), the new video coding standard released by ISO/IEC & ITU-T has just been released in its version 1 format in January 2013. It brings an additional 2:1 compression efficiency over it predecessor Advanced Video Codec (AVC) and incorporates several improvements suitable for video deployments in this new environment. This paper will examine the new improvements of HEVC, including compression performance, and what areas it may be employed to enhance in this new and changing service operator environment. Lastly this paper will conclude with integration/migration strategies to introduce HEVC technologies into Cable services, including TV Everywhere services.

INTRODUCTION

Video compression technology coupled with MPEG standardization brought a new era in video delivery and applications. The deployment of MPEG-2 video in broadcast video has been a huge success in terms of improved video quality and increased number of channels. It also increased consumer choice in the area of video services due to emergence of direct broadcast satellite (DBS), DVD and IP delivery. The MPEG-2 standard is broadcast centric and not friendly to other video applications, especially to real-time internet delivery. As the internet is a public network, and is based on a best efforts delivery protocol, streaming video delivery over the internet lacks in video quality and resolution compared to broadcast video. But, there are two primary advantages of HTTP internet delivery over broadcast delivery: 1) video/audio content can be delivered to any receiving device with internet connectivity and 2) content can be delivered in a personalized manner. To help reduce the bandwidth needed for video delivery and to address a wider area of video applications, the Joint Video Team (JVT) of MPEG and ITU-T published AVC/H.264 video compression standard in 2003 [9]. AVC provides 2:1 compression gain over MPEG-2 video. This acted as a catalyst for explosive growth in video applications, especially video over the internet. Although the internet was generally developed as a non-real-time data delivery network, numerous video applications, real-time and non-real-time, are now using nearly 50% of the internet bandwidth capacity.

Obviously this is impacting other services delivered over the internet due to real time bandwidth delivery demands. To mitigate the negative impact from video delivery over the internet, MPEG and ITU-T formed a Joint Collaboration Team for Video Coding (JCT-VC) and initiated another compression standard known as HEVC/H.265 in 2010 which provides significantly better compression than that of AVC. After working nearly 2-1/2 years, JCT-VC finalized HEVC version 1 standard in January 2013 [1, 2]. Again HEVC provides nearly 2:1 compression gain over AVC.

HEVC: DIFFERENCES FROM AVC

The development of HEVC started approximately a decade after AVC was started, but essentially is still an evolution of AVC with enhancement and refinement to some AVC tools, and with the addition of a few new tools (See Figure 1).

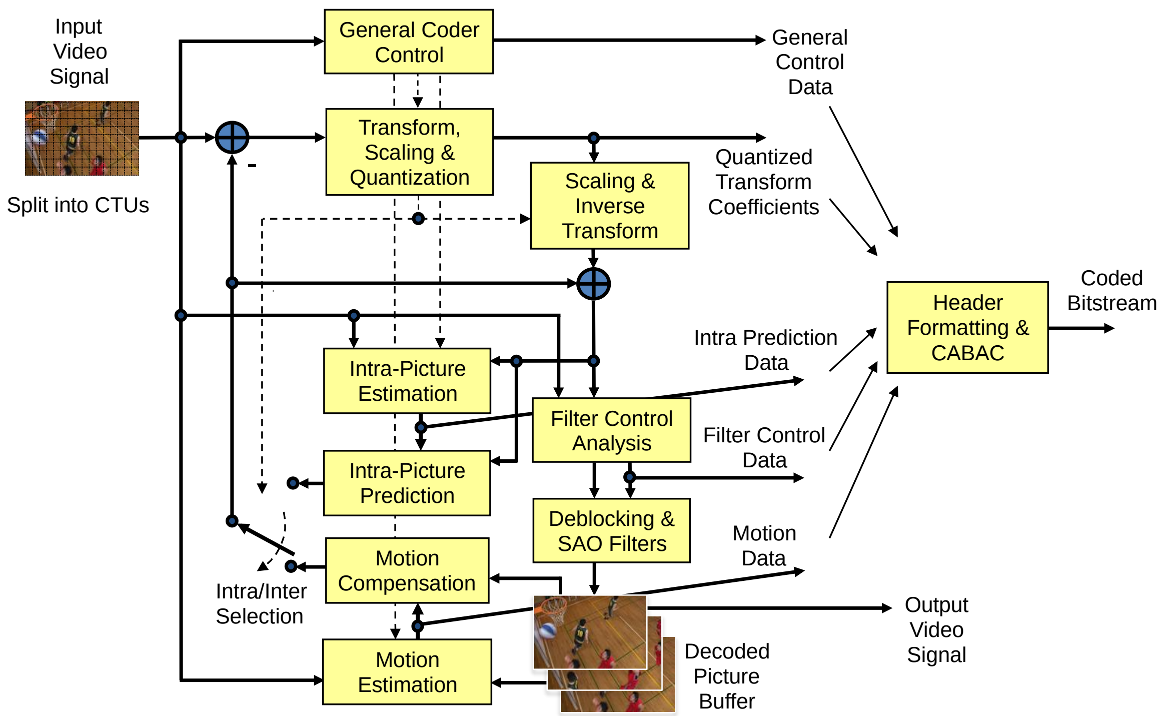

Figure 1: Typical MPEG HEVC Video Encoder Structure [10]

The new HEVC compression tools can be categorized around three main areas while maintaining the same or better visual quality by providing: 1) improvements to reduce number of bits required for region representations (Coding Unit, Transform Unit), 2) improvements for better prediction accuracy and reduction of errored residuals (e.g., Prediction Units, Spatial directional Modes, Adaptive Quantization), and 3) improvements in informational compaction/ symbol rates in bitstreams (simplified CABAC, New Scanning Modes). Table 1 below describes some of the evolution in common encoding tools from MPEG-2 to AVC to HEVC.

Coding Tools	MPEG-2	AVC	HEVC
Intra-prediction	None	Yes (9 modes)	Yes (35 predictions)
Inter-prediction	Yes (No B-picture as reference)	Yes (allows hierarchical b-picture as reference)	Same as AVC
CU (coding unit) size	16x16 (fixed, known as Macroblock (MB))	16x16 MB (same as in MPEG-2 video)	Variable, 64x64, 32x32, 16x16, and 8x8
PU (prediction unit) size	16x16	16x16, 16x8, 8x16	32x32, 16x16, 16x8, 8x16 8x8, 8x4, 4x8, 4x4
TU (transform unit size)	8x8 (DCT floating point)	8x8 and 4x4 (DCT integer)	32x32, 16x16, 8x8, 4x4 (DCT integer and also 4x4 DST integer)
In-loop filter	None	One Deblocking filter	Two in-loop filters (deblocking and SAO)
Entropy	VLC	CAVLC and CABAC	CABAC only
Parallel Processing tool	None	None	Tile and Wavefront

Table 1: Evolution in Common Video Encoding Tools

It is to be noted that the coding unit (CU) is analogous to the AVC macroblock but in this case the macroblock can change in size (see Figure 2) which when used appropriately can lead to bitrate savings. (See Table 2 [8].) Similarly, compression algorithms adaptively determine the size of prediction unit (PU) and transformation unit (TU) to achieve savings in bits while the picture quality is maintained at a desired level.

To aid in quality improvements, AVC uses a Deblocking Filter as an In-loop filter; HEVC uses a simpler but comparable Deblocking Filter and also adds a new Sample Adaptive Offset (SAO) Filter as the In-loop Filters. For Inter-picture prediction, HEVC uses Quarter-sample precision for the motion-vectors, and 7-tap or 8-tap filters for interpolation of fractional-sample positions. Whereas, AVC uses 6-tap filtering of half-sample positions followed by linear interpolation for quarter-sample positions.

HEVC also has three new features (Tiles, Wavefront Parallel Processing, and Dependent Slice Segments) to enhance parallel processing capabilities or modify slice data structures for packetization purposes. Such features help in an encoder or decoder implementation to derive benefits in particular application contexts.

Lastly HEVC also provides enhanced High Level Syntax bitstream support to improve operations over a variety of applications, network environments and robustness to data losses.

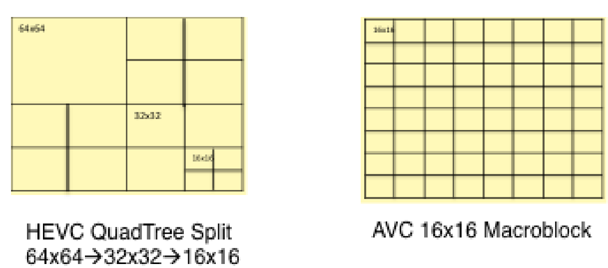

Figure 2: MPEG HEVC Coding Unit (CU) compared to MPEG AVC Macroblock

	Entertainment Applications		Interactive Applications	
	Maximum CU Size		Maximum CU Size	
	32 x 32	16 x 16	32 x 32	16 x 16
Class A	5.7%	28.2%	–	–
Class B	3.7%	18.4%	4.0%	19.2%
Class C	1.8%	8.5%	2.5%	10.3%
Class D	0.8%	4.2%	1.3%	5.7%
Class E	–	–	7.9%	39.2%
Overall	2.2%	11.0%	3.7%	17.4%
Enc. Time	82%	58%	83%	58%
Dec. Time	111%	160%	113%	161%

Table 2: Example of bitrate increase from 64x64 CU size to 32x32/16x16 CU [7]

In brief, enhancements/refinements have been done to some AVC tools with addition of a few new tools as well. The effect of these improvements has been the ability to use existing AVC tools in combinations to more precisely allocate bits in alignment with visual perceptual models as indicated by Figure 3 [6]. This allows for future advances in compression efficiency as image analysis algorithms improve. In addition, the reduction in bitrate and new parallel processing tools align well with battery-enabled portable devices. This is done through the use of low-power multi-processors CPUs rather than a single high-power processor and helps in new models of how video is consumed (connectedness, conferencing, portability, and on-demand).

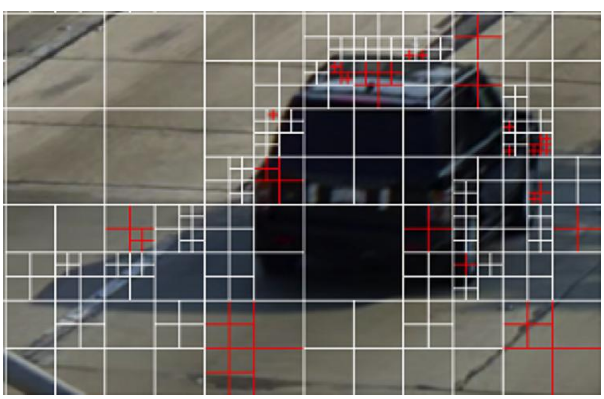

Figure 3: Example of recursive quad-tree partitioning for coding block (white) and transform block (red) [3]

PROFILES AND LEVELS

Profiles, tiers and levels specify conformance points for implementing the standard in an interoperable way. A profile defines a set of coding tools or algorithms that can be used in generating a conforming bitstream, whereas a level places constraints on certain key parameters of the bitstream such as maximum picture size, maximum bit rate, and a few other parameters which basically relate to decoder processing load and memory capabilities. In the design of HEVC, it was determined that two distinct sets of applications exist that have requirements that differ only in terms of maximum bit rate and CPB capacities. To resolve this issue, two tiers were specified for some levels – a “Main” tier is intended for most consumer video applications and a “High” tier for higher quality delivery that requires much higher bitrates. A level using the “Main” tier only needs to encompass levels below with bitrates targeted for the “Main” tier at that level. The decoders conforming to a specific profile must support all features in that profile.

Currently, three profiles, called the “Main”, “Main 10” and “Main Still” have been specified. Main is intended to be used with video with pixel depth of 8 bits and Main 10 is for 10 bit video. Main and Main 10 have the same set of tools, and a Main 10 decoder can decode a Main 8 compliant bitstream, however, the reverse is not true. It is possible that some other profiles of the standard will also be specified. Minimizing the number of profiles provides a maximum amount of interoperability among devices, and across applications such as broadcast, mobile, conferencing and streaming. Thus, a Main profile compliant device may be used for more than one application. Main Still is intended to be used for still picture decoding.

The objective here is that devices like cameras, smart phones and other similar devices can capture or decode both video as well as still pictures without the need to support two different codecs as is the case today.

The definition of 13 picture levels has been defined, starting with picture sizes such as a luma picture size of 176×144 (sometimes called quarter common intermediate format) to picture sizes as large as 7680×4320 (often called 8k). There are two tiers supported for eight of these levels (levels 4 and higher); Main tier and High tier. A decoder supporting High tier will be able to decode Main tier compliant streams, while the reverse is not true. High tier decoders support bitstream with higher maximum bitrate and hence may need higher processing power. In most cases high bitrate bitstreams provide higher video quality than the one provided by a lower bitrate bitstream.

PERFORMANCE IMPROVEMENTS

It is noted earlier that HEVC provides 2:1 better compression gain over AVC. So to deliver content with same quality, a HEVC codec will use approximate half the bitrate of that required by AVC. In other words with bitrate close to the AVC bitrate, a HEVC coded video can provide significantly better video quality. Again the compression gain can vary with content and also picture sizes. Larger size pictures tend to provide better compression efficiency than smaller resolution pictures. The table below shows the typical bitrates needed by MPEG-2, AVC and HEVC bitstreams. In subjective tests results with equivalent reproduction qualities for codecs, Table 3 shows HEVC encoders use approximately 50% less bit rate on average than AVC encoders. HEVC design is especially effective for low bit rates, high-resolution video content, and low-delay communication applications.

Resolution/FPS	HEVC	AVC	MPEG-2
UHDTV- 4K/60	8-15 Mbps	18-22 Mbps	High
HD- 1080/720P60	1.5-3.5 Mbps	5-9 Mbps	9-15 Mbps
HD- 720p30	0.8-2.0 Mbps	1.5-4 Mbps	3-5 Mbps
SD	0.4-0.7 Mbps	0.7-1.5 Mbps	2-3 Mbps

Table 3: MPEG HEVC Performance over Typical and Anticipated Video Services (Linear/VOD)

Performance Stats	HEVC over AVC [High Profile]
Encoder Complexity	~10x (or less)/Early 4x
Decoder Complexity	~1.4x
Memory	~1.25x
Memory Bandwidth	~1.25x

Table 4: MPEG HEVC Anticipated Complexity, Memory, and Memory Bandwidth Performance

Real-time software decoding [3, 4, 5] of HEVC bitstreams is very feasible on current generation devices — 1080p60 decoding on laptops or desktops, and 480p30 decoding on mobile devices.

As an example-1, HEVC software decoding of 480p to 1080p at 25 or 30 fps is possible with a single core of an ARM processor. Here the player application itself can be multi-threaded, with separate decoding and display threads (OpenGL to display video in real time). Also scaling to fit screen size and YUV to RGB color conversion can be done on the GPU during shading.

As an example-2, the software playback of 4K sequences at 60 fps is possible on a laptop where bitstreams encoded with random access main profile at 12 Mbps bit-rate can be decoded by using three parallel decoding threads with a quad core, 2.7 GHz, Core-i7 processor. Up to 100 fps are achieved with four parallel decoding threads on the same laptop.

MPEG HEVC encoders are expected to be a few to several times more complex than AVC encoders, and a subject of research in years to come. The market availability of real-time 1080@60P encoders supporting HEVC Main/Main 10 profile is expected to be available by 2014.

Service Providers: WHAT's ALREADY CHANGING?

Over the last few years, the types of services being supported have significantly evolved (see Figure 4). Previously customers watched video services through a cable set-top box and on one of the 2-3 TVs in the house. Data connections were meant for the PC and laptops available in the house. With the advent of wireless systems, those PCs, laptops, and now phone devices became portable, but bandwidth reliability was still causing interruptions and used mainly for data services rather than video viewing. With higher performing data devices, and the advent of MPEG-AVC, and settling some of the firewall issues, these devices are starting to favor video viewing more often. At the same time, television screens are getting bigger, HD is more firmly entrenched in households, and various flavors of non-linear viewing (on-demand / DVR) are increasing in popularity. With these establishing trends, the service expectation is evolving to view video on a plethora of devices (cable box, gaming console, Blu-Ray Player, Boxee Box, Roku, Tivo, Tablet, and Smartphone) and not just something that is permanently fixed in the living room. The choices made by the customer then are determined by the perceived value of the service.

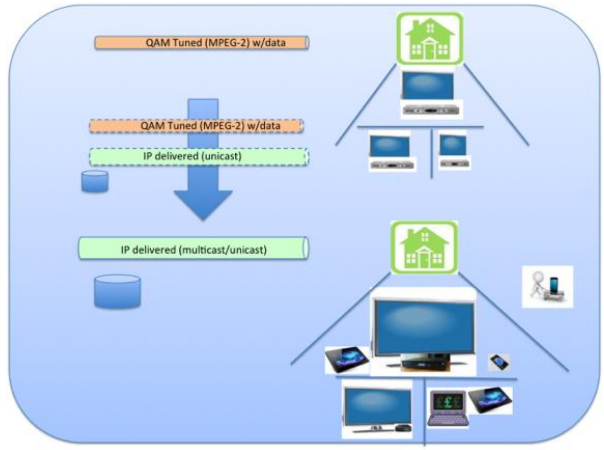

Figure 4: Household Past and Present/ Future Video Device Consumption Models

So what makes content more valued to the customer? Is it the material? Partially. But it also has a lot to do with connection and experience as well (see Figure 5). In terms of content, the value can be affected by how it can be purchased (subscription, single time, catalog depth), who is excluded from getting it, and by when it can be available (first to air, first available). In terms of experience, value can be affected by how immersive it is (large screen/4k/8k/portable screen), where it can be watched (fixed, mobile), how it can be shared (group/individual setting), and lastly how it can be consumed (languages, CC, Binge viewing). There are some exciting things happening in the area of experience. The content can be more valued because it is larger (more resolution), faster (more frame rate), and brighter (more contrast). Personal/Portable devices (PC/tablets) in the home are rapidly replacing the 3rd or 4th TV set in the household. The value of content within the household increases in this mode due to added accessibility of the content and transferability of it (your child can now watch anime on his/her laptop). The value further increases if one can then walk out the door and still can watch the content. The tradeoff on this is what the risk is to secure the content versus mobility, and a compromise to this would perhaps be a download-to-go application. In addition to increased resolution there is another dimension of increased frame-rates (60 Fps, 30 Fps and 3:2 pull-down). Lastly with advancements in cameras, images can be captured with higher contrast due to improvements in sensors which leads to shooting of a lot of dark, high-contrast scenes. In this case the experience has expanded, but is only allowed if the implemented technology can support it.

The last component is connection (which is more hidden to the service) that deals with more infrastructure and technology aspects. The value here can be affected by the availability of the connection (always-on, managed, and unmanaged), the distribution of the content (broadcast, unicast, DRM, and link encryption), the host processor capabilities (STBs, tablets, software based, hardware based), and lastly by request (linear, non-linear). Note the advent of IP delivery has opened up a set of new customer owned and managed devices. The choices of using broadcast or IP unicast can limit or expand the audience size when demand requires it. In earlier periods, the infrastructure design deemed all content valuable and protected the content all equally, but restricted content to the confines of the infrastructure. With the advent of DRM-protected content, the infrastructure restrictions were loosened, all content did not have to be protected equally, and it became more mobile. Lastly technical capabilities, like codec format, can affect capacity of the connection because the codec can affect things like the management of bandwidth, the capacity of storage, and the processing on devices. The calculations are for the value of content changes to include platform portability as well as increasing the weighting of accessibility and availability. In the real-time linear domain, this can increase the value of sports and news. While in the VOD domain, this can help determine if a cable subscription of aggregate VOD catalog is ordered, or whether an over-the-top (OTT) service account is enough. What is also interesting is that the infrastructure and technology can change but the service to the customer could remain the same.

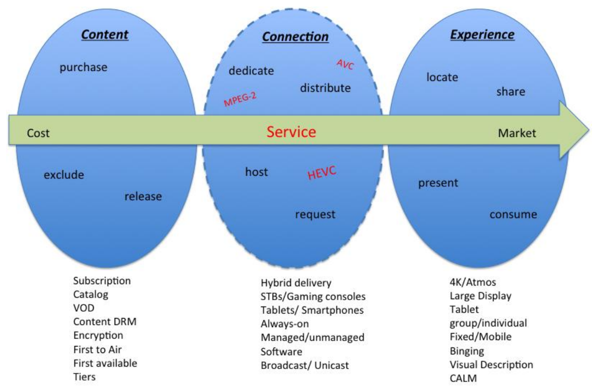

Figure 5: Content & Service Value Diagram

The service evolution that has been happening has affected the connection value aspects the most. With the new devices and new distribution pathways to get content in front of the viewer, there has been a subsequent 10x-100x increase of the number of streams and files to support new devices and unmanaged networks in the service operator ecosystem. These changes have put increased pressure on bandwidth and storage demands in the network due to support of adaptive streaming, expanded on-demand catalogs, and PVR services. Until just recently a single transcoder would have a single output stream. Presently a single input could generate up to 10 streams or files at the output. New distribution pathways such as on-demand MBR (Multi-BitRate) streaming for http adaptive streaming technologies has also increased bandwidth and storage demands to support bandwidth OTT congestion management and network storage.

The delivery of video is not only to a STB using an MPEG-2 transport stream over QAM, but now also to connected personal devices through an HTTP protocol over IP. For individual viewing of real-time content, the IP connection is very suitable to this medium. But when it comes to high valued content that is driven real-time high viewer demanded, often the MPEG-TS system has benefits that are sometimes more complicated to replicate in the IP domain. To repeat this experience in the IP domain would be an analogous multicast equivalent similar to the broadcast medium, but in the near term it is still replicated using a unicast IP delivery. Presently the storage may be commonly shared (used in HTTP adaptive streaming) but the delivery is still a unicast delivery which works unless one scales up in real-time viewership (e.g., the Super Bowl).

The service expectation is already changing to support better displays, more mobile and personal devices, and increased non-linear viewing. The service expectation has been moving in this direction because of the increase in value to the customer experience in the living room, and on the go. A large amount of improvement and value to the services has already happened with more expected to come. There will be increased pressure on the infrastructure due to higher pixel demands, increased number of output streams, and related storage requirements. Can our current technologies (Bandwidth, Storage, IP distribution) support the ramp up of these services as popularity grows? HEVC can be a good candidate to relieve some of the infrastructure demands as the service continues to grow.

POSSIBLE MPEG HEVC BENEFITS

HEVC can reduce the complexity and costs of handling multiple streams in this transitioning environment. For OTT-based services, the number of stream representations can be reduced since higher quality streams at lower bitrates can survive more often through bandwidth-congested environments. For a present HD OTT service, the number of streams could be reduced by 50% and in addition each of the remaining streams has also a reduced bitrate. For backhaul, bandwidth distribution demands for mezzanine/contribution streams can be reduced by switching from AVC/MPEG-2 formats to HEVC (or the quality of distributed content can be increased for the same bandwidth costs). This also encourages the transition from satellite-based feeds to fiber-based IP connections, even further reducing the signal integrity risks encountered with satellite distribution and redundancy strategies associated with it. In the area of non-linear services (VOD/Cloud DVR) and targeted ad insertion services, storage demands can be reduced by storing in an HEVC format even though output may be transcoded to a more traditional format. This becomes especially needed in light of supporting unique copy services across cable devices and customer owned and managed (COAM) devices. The reduction of storage demands can be greater than 50% due the decrease in number of streams supported and the decrease in bitrate to support the same quality streams. At the granularity of the stream for VOD, efficiency can occur by switching from using multiple trick files to support 2-3 speeds to supporting dynamic trickplay in the stream itself through the use of picture order count (poc) / temporal IDs in HEVC.

In terms of the device perspective, HEVC can increase the value associated with the host device by decreasing the service bandwidth to each device and increasing the amount of bandwidth available for the service. This benefit will assist in terms of IP unicast replications of expected broadcast customer experiences (e.g., popular events like the Super Bowl) while longer term solutions can then be created (e.g., multicast IP). Additionally HEVC does not require a new set of mobile personal devices to be created; existing mobile devices can already play 720p30 HEVC using software-based decoders taking advantage of the small increase in decoder complexity. Since mobile devices often have short (18 month) lifetimes, the transition to HEVC in this area can be easier with later introduction of higher valued content experiences. 720p30 HD streams are already the currently accepted deployment of HD streams for OTT services and smartphone/tablets devices. Lastly, battery lifetimes on the device (often consumed quickly by video viewing) can increase (HEVC streams with 4-8 hours of continuous play) by switching to an HEVC coding standard designed to take advantage of low-power multi-core processors in these devices. Additionally the reduced bitrate can save on antennae power and processing power required to decode on the portable device. This can also further increase battery lifetime.

In the areas with more immersive experiences and larger displays such as in the living room/bedroom, HEVC can greatly help out in this area by reducing the bitrate associated with carrying a higher quality or larger resolution stream. In terms of integrating these new experience streams, HEVC can reduce the bit rate to fit into the same slot as a present day MPEG-2 HD channel (~8-15 Mbps). With higher resolution streams, the gains could increase non-linearly (compared to HD) due to use of larger macroblock sizes, and increased spatial mode. With increased frame rates, HEVC can also reduce bit rate from adding additional frames using simplification methods/increasing accuracy in motion vectors and use of different options in prediction units. Hardware based decoders through silicon based chipsets are coming out early next year with 4K decoding that will add value to the larger display based devices that can be implemented through an ROI based approach based on 4K and higher services. Bandwidth demands on the plant can be managed with expectations set for new higher quality deployments (4K@60fps, 10 bit displays, etc.) within the same bandwidth of a traditional HD and in combination with other bandwidth saving strategies. The install base on this would coincide with an increased value in services that can be enabled on next generation STBs and gaming machines.

CONCLUSIONS

Cable services have evolved rapidly from a few years ago. Video services are now being offered for not just the living room, but also on the go for laptops, tablets, and mobile devices. An IP distribution structure using adaptive streaming techniques is used for these newer devices which enables portability in the managed home environment, WIFI, and OTT) delivery. Additionally there is an increase in number of streams delivered and managed as well as storage in cloud DVR and CDN structures. With the increased popularity of these services, there will be increasing pressure on the current infrastructure to support these services. Higher Efficient Video Coding (HEVC) can play an important role in deploying these new services on a larger scale by reducing the number and size of streams delivered and stored, making the IP unicast model more efficient for replication of older broadcast services, increasing the battery lifetimes on devices, and creating more immersive experiences for larger displays with higher frame rates and resolutions. HEVC can increase value by making the connection more efficient and improving the experience which can support the expansion of the service market.

ENDNOTES

[1] B. Bross, et. al., “High Efficiency Video Coding (HEVC) text specification draft 10 (for FDIS & Last Call)”, JCTVC-L1003_v34, Geneva, January 2013

[2] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, “Overview of the High Efficiency Video Coding (HEVC) Standard,” IEEE Transactions on Circuits and Systems for Video Technology, December, 2012.

[3] F. Bossen, B. Bross, K. Suhring and D. Flynn, “HEVC Complexity and Implementation Analysis”, IEEE Transactions on Circuits and Systems for Video Technology, December, 2012.

[4] T.K. Tan, Y. Suzuki and F. Bossen, “On software complexity: decoding 4K60p content on a laptop”, JCTVC-L0098, Geneva, CH, 14–23 January 2013.

[5] K. Veera, R. Ganguly, B. Zhou, N. Kamath, S. Chowdary, J. Du, I.-S. Chong, and M. Coban, “A real-time ARM HEVC decoder implementation,” JCTVC- H0693, San Jose, CA, February 2012.

[6] Y. Syed and D. Holden, “Strategies for Deploying High Resolution and High Frame Rate Cable Content Leveraging Visual Systems Optimizations,” Proceeding of NCTA 2012, May 21, 2012.

[7] J.-R. Ohm, G. Sullivan, H. Schwarz, T.K. Tan and T. Wiegand, “Comparison of the Coding Efficiency of Video Coding Standards—Including High Efficiency Video Coding (HEVC)”, IEEE Transactions on Circuits and Systems for Video Technology, December, 2012

[8] A. Rodriguez and K. Morse, “HEVC-Driving Disruption in Multiscreen Converged Service Delivery Architecture,” Proceedings of 2013 NAB Broadcast Engineering Conference, April 7, 2013, pp 33-38.

[9] ITU-T Rec. H.264 | ISO/IEC 14496-10, (2005), “Information Technology – Coding of audio visual objects – Part 10: Advanced Video Coding.”

[10] Gary J. Sullivan, HEVC Tutorial in INCITS meeting at CableLabs, March 2013.

HIPnet: A self-configuring multi-router Home IP network.

Chris Grundemann
CableLabs

Abstract

There are many new pressures and requirements emerging in today's home networks: The need for separation of visiting guest users from home users, community Wi-Fi services, smart grid, home automation & security, and an ever increasing number and type of IP enabled devices in the subscriber home are all strong motivations for additional routers and multiple LANs. The emergence of heterogeneous link layer technologies, machine to machine communication, IP & multicast video streaming, video content sharing, telecommuting and corporate IT requirements, and the possibility of home network multi-homing are all also driving additional complexity and new requirements into home networks.

This paper presents a novel approach to home router architecture, which applies many of the tools and protocols within the IPv6 framework in new ways in order to enable a completely self-configuring dual-stack (IPv4 & IPv6) multi-router home network capable of supporting the full range of in-home IP services. While many in this field are focusing on routing protocols and other complex, long-term solutions, the HIPnet approach leverages the existing Neighbor Discovery (ND) and DHCPv6 protocols, making it simpler and cheaper to implement in the near term while being robust enough to work for the long-term as well.

The paper explains the idea of directionless home routers, which have no hard-set LAN or WAN ports but rather use our "up detection" mechanism to elect a WAN port from the available physical interfaces in a deterministic manner. It describes how this method of up detection is able to create a logically hierarchical network even in a completely arbitrary and loop-filled physical topology, without introducing any new protocols. This paper then introduces "recursive DHCP Prefix Delegation (PD)" and an algorithm for IPv6 prefix sub-delegation, where the prefix delivered to the home router is divided into smaller sub-prefixes that are then distributed to directly connected downstream routers. It also explains a method for using bits from those delegated IPv6 prefixes to seed unique IPv4 prefixes without the need for IPv4 prefix delegation in DHCP. Finally, the paper describes hierarchical routing and how this entire system works as a whole to enable ISP failover and limited multihoming.

INTRODUCTION

Yesterday's home network

Home networks of the past have largely fit the same basic model: One home router connecting one Internet Service Provider (ISP) to one in-home Local Area Network (LAN). The IPv4 home LAN typically uses [RFC1918] "private" IPv4 address space to number networked devices. These "private" addresses are then rewritten in the IP header using Network Address Translation (NAT) overloading at the home router to allow a single "public" (globally routable) IPv4 address to be shared among all IP enabled devices in the home netowork for Internet connectivity.

More recently, many home users have started to add additional routers to their home networks, often unintentionally. It is becoming common to find home users who has purchased and deployed a second or third home router in order to extend Wi-Fi range or provide additional Ethernet ports. Although there are devices better suited to these tasks (Access Points (APs) and Repeators for Wi-Fi range and Ethernet Switches for physical ports), these "specialty" devices are less common in retail stores and often cost as much or more than the general purpose routers. Additionally, many people (consumers and retail store clerks alike) are more familiar with home routers than they are other home networking gear. This combination of familiarity, availability, and affordability seems to be driving more and more home users to (often inadvertently) deploy multi-router home network topologies.

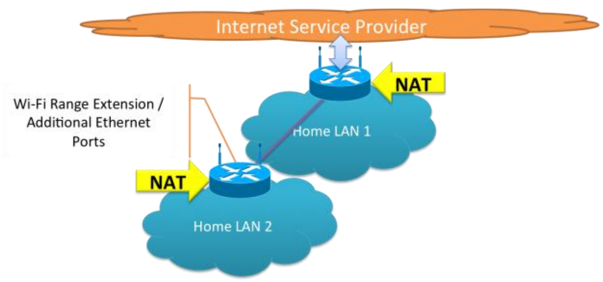

Yesterday's Home Network

In the legacy paradigm of IPv4-only and NAT, this deployment of multi-router home networks is problematic. Because these routers are designed with the one-router/one-LAN architecture in mind, there are several problems introduced when using more than one of them in a home network.

First, adding an additional router creates an additional LAN, with it's own DHCP server, address pool (which often overlaps other home LANs), and default route (through the new router). As this is usually not the intention of the home user, and may not even be a known consequence, it can cause problems when setting up new devices and services or when troubleshooting existing ones. Introducing multiple, routed LANs also stops link-local traffic from reaching the entire home network, which breaks many forms of service discovery.

Second, inserting even a second legacy home router creates introduces a second NAT, this one within the home network itself. This causes traffic between the home LANs to undergo NAT and all traffic leaving the second LAN to undergo two rounds of NAT before reaching the ISP network. These in-home and multi-layer NATs are known to impair or completely break many protocols and applications (e.g. service discovery, IPsec, DNSSEC, etc.). Additional layers of routers add additional layers of NAT, worsening the problem.

Third, these NAT routers inherently introduce a stateful firewall, which exaserbates many of the issues already raised.

Perhaps even more alarming is the state of IPv6 in the home. Unlike IPv4 NAT which allows multiple routers to be linked one behind the other to offer at least some connectivity to connected devices, current IPv6 enabled home routers [6204bis] do not support any standard mechanism to facilitate such "chaining." This means that devices connected to a second home router are likely to not have any IPv6 connectivity outside of the LAN at all.

The end result is that multi-router home networks built with legacy home routers have limited functionality, and the functionality is further limited as size and complexity increase. There are ways to solve many of these limitations but they all require manual configuration of home routers and networked devices, which is beyond the ability of most home users.

<u>Emerging use cases</u>

While home IP networks have been able to remain relatively simple over their 20-30 year life so far, there are now use-cases emerging which will surely change that going forward.

One of the first trends to appear is that of "guest" networks. Many home users have an increasing amount of personal or private data on their networked devices which they may not want visitors to their home to be able to access. Family photographs, financial and tax documents, other legal materials, and many other common types of electronicly stored information are seen as sensitive. In order to provide Internet access for guests without exposing this sensitive data home users are beginning to deploy a second, "guest," SSID on their wireless networks.

Other requirements for additional home LANs include:

1) Community Wi-Fi applications where a Wi-Fi gateway in the user's home is used to provide Wi-Fi roaming services for subscribers other than the home user.

2) Femto cell applications in which a gateway in the subscriber home is used to provide cellular services.

3) Telecommuting and corporate IT requirements for network separation between business and personal LANs within the user's home.

4) The emergence of heterogeneous link layer technologies, such as Bluetooth, ZigBee, and Z-Wave, which require their own proprietary gateways.

5) The introduction of IP based or otherwise networked Security, Monitoring, and Automation systems and services.

In addition to these emerging use-cases for multi-LAN/multi-router home networks, there are several other trends towards more complex home networks. These primarily revolve around the growing amount of IP video and the ever increasing number of IP enabled devices in the subscriber home.

IP is becoming a prominent medium for transmitting video. Many home users are already streaming at least some portion of the video they consume over IP from the Internet. Likewise, IP is the de-facto standard for video content sharing and streaming between devices inside the home. This IP delivery of video places new burdens and requirements on home networks, which drives additional complexity.

The explosion of IP enabled devices in each home shows no signs of slowing in the forseable future. The smorgasboard of desktops, laptops, tablets, and mobile phones is being supplemented further with all manner of "smart" appliances, cameras, sensors, printers, storage devices, and surely more to come. This trend adds to those above and drives even more complexity into home networks.

<u>The home network of tomorrow</u>

The emerging home network use-cases outlined above, in addition with others not covered here and perhaps not even apperent yet, lead to a future in which complex, multi-router, home networks become commonplace.

A typical home network of tomorrow is likely to have some combination of multiple W-Fi SSIDs, a "guest" LAN or two, multiple function-specific LANs with their own routers or gateways (e.g. a LAN in the kitchen for smart appliances and a LAN in the living room for media devices, etc.), and a multitude of IP enabled services and devices running over it all.

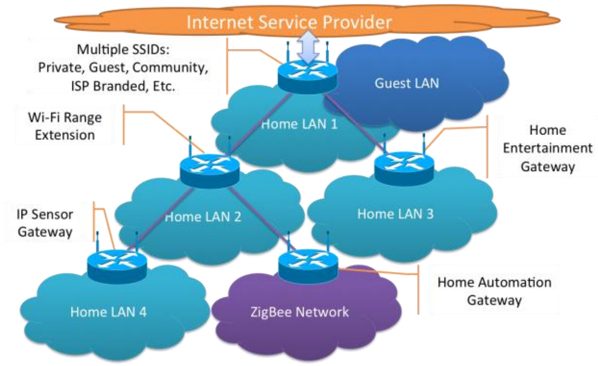

Tomorrow's Home Network

While we can not assume to predict the future with absolute certainty, we can easily see that home networks will become more complex, continuing to support more and more routers, devices, and serivces. We can also assume that these complex, multi-router home networks will need to "just work" in order to facilitate their operation by average home users.

THE HIPNET SOLUTION

HIPnet (derived from Home IP networking) is a near term solution to complex home networks. Specifically, the HIPnet solution defines a self-configuring home router architecture which:

1) Is capable of operating in increasingly large (and arbitrarily constructed) residential home networks.

2) Requires no user interaction for the vast majority of use-cases.

3) Uses existing protocols in new ways.

4) Does not require a routing protocol.

5) Meets the principles for advanced home networks defined in [homenet].

Guiding principles

Five common principles have guided the development of HIPnet:

1) *Home networks will become more complex, home users will not.* As discussed above, a multitude of emerging use-cases are driving additional complexity into home networks. Despite this, there is no reason to assume a parallel trend of home users becoming more networking savvy will emerge as well. These complex home networks of the future need to "just work" in the majority of cases for the majority of users, without any manual configuration whatsoever.

2) *Invoking a "god box" leads to religious wars.* A "god box" commonly refers to any device which attempts to be all things to all people. Here specifically a "god box" would be a home gateway intended to fulfill all home networking requirements. This is unlikely to be successful short-term in an environment of proprietary solutions nor long-term in a field that evolves as quickly as networking. Competing standards and solutions are sure to "war" with each other and make consensus on any comprehensive all-in-one solution nearly or totally impossible.

3) *New protocols bring new problems.* The introduction of any new protocol, whether completely new or simply new to that application or use, almost invariably creates new problems which must be addressed or worked around. It is almost always preferable to use the protocols already available, perhaps in slightly new ways, then to introduce the uncertainty of a new protocol.

4) *We have enough addresses.* IPv6 provides a glut of individual addresses. This allows us to be less concerned with the highest possible efficiency in address usage and to focus on simplifying network functionality.

5) *Use IPv6, support IPv4.* The IPv6 protocol suite provides us with all the tools we need to create a complete solution to complex multi-router home networks. When

creating this IPv6 based architecture, however, we mustn't forget to maintain support for legacy IPv4 devices and services.

Terminology

The following terms will be used throughout the remainder of this document to describe the HIPnet solution in detail:

Home IP Network (HIPnet) Router: A node intended for home or small-office use that forwards packets not explicitly addressed to itself.

End-User Network: One or more links attached to the HIPnet router that connect IPv6 and IPv4 hosts.

Service provider: An entity that provides access to the Internet. In this document, a service provider specifically offers Internet access using IPv6, and may also offer IPv4 Internet access. The service provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.

Customer Edge Router (CER): A HIPnet router that connects the end-user network to a service provider network.

Internal Router (IR): An additional HIPnet router deployed in the home or small-office network that is not attached to a service provider network. Note that this is a functional role; it is expected that there will not be a difference in hardware or software between a CER and IR, except in such cases when a CER has a dedicated non-Ethernet WAN interface (e.g. DSL/cable/ LTE modem) that would preclude it from operating as an IR.

Up interface: A HIPnet router's attachment to a link where it receives one or more IP addresses and/or prefixes. This is also the link to which the HIPnet router points its default route.

Down interface: A HIPnet router's attachment to a link in the end-user network on which it distributes addresses and/or prefixes. Examples are Ethernet (simple or bridged), 802.11 wireless, or other LAN technologies. A HIPnet router may have one or more network-layer down interfaces.

Downstream router: A router directly connected to a HIPnet router's Down Interface.

Depth: The number of layers of routers in a network. A single router network would have a depth of 1, while a router behind a router behind a router would have a depth of 3.

Width: The number of routers that can be directly subtended to an upstream router. A network with three directly attached routers behind the CER would have a width of 3.

Edge detection

Customer Edge Routers (CER) will often be required to behave differently from Internal Routers (IR) in several capacities. Some examples include: Firewall settings, IPv4 NAT, ULA generation (if supported), name services, multicast forwarding differences, and others. This is a functional role, and will not typically be differentiated by hardware/software (i.e. end users will not purchase a specific CER model of router distinct from IR models).

There are three methods that a router can use to determine if it is a CER for its given network:

1) "/48 Check" - Service providers will provide IPv6 WAN addresses (DHCPv6 IA_NA) and IPv6 prefixes (DHCPv6 IA_PD) from different pools of addresses. The largest IPv6 prefix that we can expect to be delegated

to a home router is a /48. Combining these two observations, a home router can compare the WAN address assigned to it with the prefix delegated to it to determine if it is attached directly to a service provider network. If the router is a CER, the WAN address will be from a different /48 than the prefix. If the router is an IR, the WAN address will be from the same /48 as the prefix. In this way, the router can determine if it is recieving an "external" prefix from a service provider or an "internal" prefix from another home router.

2) CER_ID - A home router can use the CER_ID DHCPv6 option defined in [CER-ID] to determine if it is a CER or an IR. ISPs will not set the CER_ID option, but the first CPE router sets its address in the option and other routers forward the completed CER_ID to subdelegated routers.

3) Physical - Some routers will have a physical differentiator built into them by design that will indicate that they are a CER. Examples include mobile routers, DSL routers, and cable eRouters. In the case of a mobile router, the presence of an active cellular connection indicates that the router is at the customer edge. Likewise, for an eRouter, the presence of an active DOCSIS® link tells the router that it is at the customer edge.

HIPnet routers can (and likely will) use more than one of the above techniques in combination to determine the edge. For example, an internal router will check for the CER_ID option, but will also use the 48 check in case its upstream router does not support CER_ID.

Directionless routers

As home networks grow in complexity and scale, it will become more common for end users to make mistakes with the physical connections between multiple routers in their home or small office. This is liklely to produce loops and improper uplink connections. While we can safely assume that home networks will become more complex over time, we cannot make the same assumption of the users of home networks. Therefor, home routers will need to mitigate these physical topology problems and create a working multi-router home network dynamically, without any end user intervention.

Legacy home routers with a physically differentiated uplink port are "directional;" they are pre-set to route from the 'LAN' or Internal ports to a single, pre-defined uplink port labeled "WAN" or "Internet". This means that an end-user can make a cabling mistake which renders the router unusable (e.g. connecting two router's uplink ports together). On the other hand, in enterprise and service provider networks, routers are "directionless;" that is to say they do not have a pre-defined 'uplink' port. While directional routers have a pre-set routing path, directionless routers are required to determine routing paths dynamically. Dynamic routing is often achieved through the implementation of a dynamic routing protocol, which all routers in a given network or network segment must support equally. This section introduces an alternative to dynamic routing protocols (such as OSPF) for creating routing paths on the fly in directionless home routers.

Note that some routers (e.g. those with a dedicated wireless/DSL/DOCSIS® WAN interface) may continue to operate as directional routers. The HIPnet mechanism described below is intended for general-purpose routers.

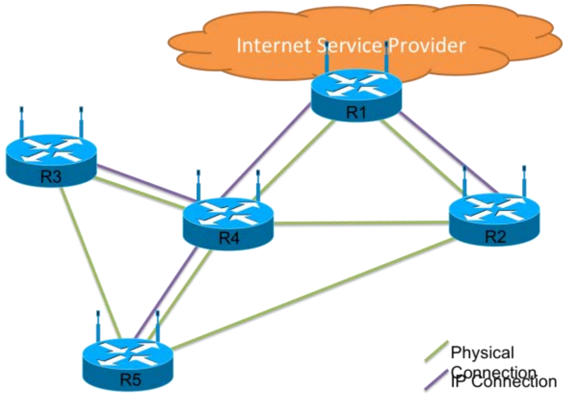

Physically Arbitrary with Logical Hierarchy

The HIPnet mechanism uses address acquisition as described in [6204bis] and various tiebreakers to determine directionality (up vs. down) and by so doing, creates a logical hierarchy (cf. [prefix-alloc]) from any arbitrary physical topology:

1) After powering on, the HIPnet router sends Router Solicitations (RS) [RFC4861] on all interfaces (except Wi-Fi*)

2) Other routers respond with Router Advertisements (RA)

3) Router adds any interface on which it receives an RA to the candidate 'up' list

4) The router initiates DHCPv6 PD on all candidate 'up' interfaces. If no RAs are received, the router generates a /48 ULA prefix.

5) The router evaluates the offers received (in order of preference):

a) Valid GUA preferred (preferred/valid lifetimes >0)

b) Internal prefix preferred over external (for failover - see below)

c) Largest prefix (e.g. /56 preferred to /60)

d) Link type/bandwidth (e.g. Ethernet vs. MoCA)

e) First response (wait 1 s after first response for additional offers)

f) Lowest numerical prefix

6) The router chooses the winning offer as its Up Interface.

Once directionality is established, the router continues to listen for RAs on all interfaces but doesn't acquire addresses on Down Interfaces. If the router initially receives only a ULA address on its Up Interface and GUA addressing becomes available on one of its Down Interfaces, it restarts the process. If the router stops receiving RAs on its Up Interface, it restarts the process.

In all cases, the router's Up Interface becomes its uplink interface; the router acts as a DHCP client on this interface. The router's remaining interfaces are Down Interfaces; it acts as a DHCP server on these interfaces. Also, per [6204bis], the router only sends RAs on Down Interfaces.

*Note: By default, Wi-Fi interfaces are considered to point "down." This requires manual configuration to enable a wireless uplink, which is preferred to avoid accidental or unwanted linking with nearby wireless networks.

<u>Recursive prefix delegation</u>

HIPnet routers use DHCPv6 prefix sub-delegation ([RFC3633]) to recursively build a hierarchical network ([prefix-alloc]). This approach requires no new protocols to be supported on any home routers.

Once directionality is established, the home router will acquire a WAN IPv6 address and an IPv6 prefix per [6204bis]. As HIPnet routers (other than the CER) do not know their specific location in the hierarchical network, all HIPnet routers use the same generic rules for recursive prefix delegation to facilitate route aggregation, multihoming, and IPv4 support (described below). This methodology expounds upon that previously described in [prefix-alloc].

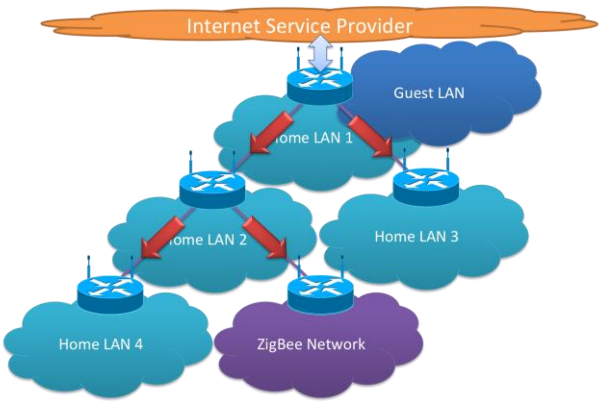

Recursive Prefix Delegation

The process can be illustrated in the following way:

1) Per [6204bis], the HIPnet router assigns a separate /64 from its delegated prefix(es) for each of its Down Interfaces in numerical order, starting from the numerically lowest.

2) If the received prefix is too small to number all Down Interfaces, the router collapses them into a single interface, assigns a single /64 to that interface, and logs an error message.

3) The HIPnet router subdivides the IPv6 prefix received via DHCPv6 ([RFC3315]) into sub-prefixes. To support a suggested depth of three routers, with as large a width as possible, it is recommended to divide the prefix on 2-, 3-, or 4-bit boundaries. If the received prefix is not large enough, it is broken into as many /64 sub-prefixes as possible and an error message is logged. By default, this document suggests that the router divide the delegated prefix based on the aggregate prefix size and the HIPnet router's number of physical Down Interfaces. This is to allow for enough prefixes to support a downstream router on each down port.

* If the received prefix is smaller than a /56 (e.g. a /60), a HIPnet router with 8 or more ports divides on 3-bit boundaries (e.g. /63) and a HIPnet router with 7 or fewer ports divides on 2-bit boundaries (e.g. /62).

* If the received prefix is a /56 or larger, a HIPnet router with 8 or more ports divides on 4-bit boundaries (e.g. /60) and a HIPnet router with 7 or fewer ports divides on 3-bit boundaries (e.g. /59).

4) The HIPnet router delegates remaining prefixes to downstream routers per [RFC3633] in reverse numerical order, starting with the numerically highest. This is to minimize the renumbering impact of enabling an inactive interface.

For example, a four port router with two LANs (two Down Interfaces) that receives 2001:db8:0:b0::/60 would start by numbering its two Down Interfaces with 2001:db8:0:b0::/64 and 2001:db8:0:b1::/64 respectively, and then begin prefix delegation by giving 2001:db8:0:bc::/62 to the first directly attached downstream router.

<u>Hierarchical routing</u>

The recursive prefix delegation method described above, coupled with "up detection", enables very simple hierarchical routing. By this we mean that each router installs a single default 'up' route and a more specific 'down' route for each prefix delegated to a downstream IR. Each of these 'down' routes simply points all packets destined to a given prefix to the WAN IP address of the router to which that prefix was delegated. This combination of a default 'up' route and more specific 'down' routes provides complete reachability within the home network with no need for any additional message exchange or routing protocol support.

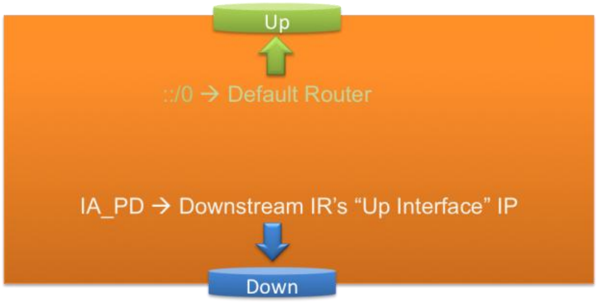

Hierarchical Routing Table

Multiple address families

The recursive prefix delegation method described above can be extended to support additional address types such as IPv4, additional GUAs, or ULAs. When the HIPnet router receives its prefix via DHCPv6 ([RFC3633]), it computes its 16-bit Link ID (bits 48-64) from the received IA_PD. It then prepends additional prefixes received in one or more IPv6 Router Advertisements ([RFC4861]) or from the DHCPv4-assigned ([RFC2131]) IPv4 network address received on the Up Interface.

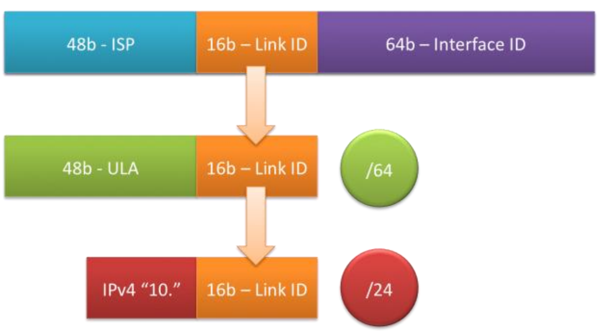

Link ID

As the network is hierarchical, upstream routers know the Link ID for each downstream router, and know the prefix(es) on each LAN segment. Accordingly, HIPnet routers automatically calculate downstream routes to all downstream routers.

In networks using this mechanism for IPv4 provisioning, it is suggested that the CER use addresses in the 10.0.0.0/8 ([RFC1918]) range for downstream interface provisioning.

Multiple ISP: Failover

Using the procedures described above, multi-router home networks with multiple ISP connections can easily operate in an active/standby manner, switching from one Internet connection to the other when the active connection fails. Lacking a default priority, HIPnet routers will have to default to a "first online" method of primary CER selection. In other words, by default, the first CER to come online becomes the primary CER and the second CER to turn on becomes the backup. In this text, the primary ISP is the ISP connected to the primary CER and the backup ISP is simply the ISP atached to the backup CER.

In an active/standby multi-ISP scenario, a backup CER sets its Up Interface to point to the primary CER, not the backup ISP. Hence, it does not acquire or advertise the backup ISP prefix. Instead, it discovers the internally advertised GUA prefix being distributed by the currently active primary CER.

In the case of a primary ISP failure, per [6204bis], the CER sends an RA advertising the preferred lifetime as 0 for the ISP-provided prefix, and its router lifetime as 0. The backup CER becomes active when it sees the primary ISP GUA prefix advertised with a preferred lifetime of 0. In the case of CER failure, if the backup CER sees the Primary CER stop sending RAs altogether, the Backup CER becomes active.

When the backup CER becomes active, it obtains and advertises its own external GUA. When advertising the GUA delegated by its ISP, the backup CER sets the valid, prefered, and router lifetimes to a value greater than 0. Other routers see this and re-determine the network topology via "up" detection, placing the new CER at the root of the new hiearchical tree.

Using this approach, manual intervention may be required to transition back to the primary ISP. This prevents flapping in the event of intermittent network failures. Another alternative is to have a user-defined priority, which would facilitate pre-emption.

Multiple ISP: Multi-homing

The HIPnet algorithm also allows for limited active/active multihoming in two cases:

1) When one ISP router is used as the primary connection and the second ISP router is used for limited connectivity e.g. for a home office.

2) When both ISP routers are connected to the same LAN segment at the top of the tree.

In case 1, the subscriber has a primary ISP connection and a secondary connection used for a limited special purpose. (e.g. for work VPN, video network, etc.). Devices connected under the secondary network router access the Internet through the secondary ISP. All devices still have access to all network resources in the home. Devices under the secondary connection can use the primary ISP if the secondary fails, but other devices do not use the secondary ISP.

As described above, the primary CER performs prefix sub-delegation to create the hierarchical tree network. The secondary edge router then obtains a second prefix from ISP2 and advertises the ISP2 prefix as part of its RA. The Secondary CER thus includes sub-prefixes from both ISPs in all IA_PD messages to downstream routers with the same "router id.". In a change from the single-homing (or backup router) case, the secondary CER points its default route to ISP2, and adds an internal /48 route to its upstream internal router (e.g. R1). Devices below the the secondary CER (e.g. Host 2, Host 3) use ISP2, but have full access to all internal devices using the ISP1 prefix (and/or ULAs). If the ISP2 link fails, the secondary CER points its default route 'up' and traffic flows to ISP1. Devices not below the secondary CER (e.g. Hosts 1, 4, 5) use ISP1, but have full access to all internal devices using the ISP1 prefix (or ULAs).

In case 2, the secondary CER is installed on the same LAN segment as the primary CER. As above, it acquires a prefix from both the CER and secondary ISP. Since it is on the same LAN segment as the CER, the secondary CER does not delegate prefixes to that interface via DHCP. However, it does generate an RA for the ISP2 prefix on the LAN.

As described above, downstream routers receiving the secondary CER RA acquire an address using SLAAC and generate a prefix for sub-delegation by prepending the secondary CER prefix with the Link ID generated during the receipt of the prefix from the CER. Such routers then generate their own RAs on downstream interfaces and include the secondary prefix as an IA_PD option in future prefix delegations.

CONCLUSION

HIPnet is a near term, scalable solution to the problem of increasing complexity in home networks. The HIPnet architecture described in this document meets the challenges of tomorrow's home networks without the need for manual configuration or routing protocols by employing directionless home routers, recursive prefix delegation, and hierarchical routing.

While the primary focus is on IPv6 support, this document also describes how HIPnet leverages IPv6 to configure IPv4 in a manner better than the nested NATs in operation on many networks today.

This document describes how a HIPnet router automatically detects both the edge of the customer network and its upstream interface, how it subdivides an IPv6 prefix to distribute to downstream routers, and how it leverages IPv6 address assignment to distribute IPv4 addresses. It also discusses how such a router can operate with a backup ISP or limited multihoming across two ISPs.

Specific requirements for building a HIPnet compliant home router can be found in [hipnet].

REFERENCES

[prefix-alloc] Nordmark, E., Chakrabarti, S., Krishnan, S., and W. Haddad, "Simple Approach to Prefix Distribution in Basic Home Networks", draft-chakrabarti-homenet-prefix-alloc-01 (work in progress), October 2011.

[cer-id] Donley, C. and C. Grundemann, "Customer Edge Router Identification Option", draft-donley-dhc-cer-id-option-01 (work in progress), September 2012.

[hipnet] Grundemann, C., Donley, C., Brzozowski, J., Howard, L., and V. Kuarsingh, "A Near Term Solution for Home IP Networking (HIPnet)", draft-grundemann-homenet-hipnet-01 (work in progress), February 2013.

[homenet] Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "Home Networking Architecture for IPv6", draft-ietf-homenet-arch-07 (work in progress), February 2013.

[6204bis] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-6204bis-12 (work in progress), October 2012.

[RFC 1918] IETF RFC 1918, Address Allocation for Private Internets, Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, February 1996.

[RFC 2131] IETF RFC 2131, Dynamic Host Configuration Protocol, R. Droms, March, 1997.

[RFC 3315] IETF RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), R. Droms, Ed., J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney, July 2003.

[RFC 3633] IETF RFC 3633, IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6, O. Troan, R. Droms, December 2003.

[RFC 4861] IETF RFC 4861, T. Narten, E. Nordmark, W. Simpson, H. Soliman, Neighbor Discovery for IP Version 6 (IPv6), September 2007.

Improving Adaptive Video Delivery Through Active Management

Santhana Chari
Arris
santhana.chari@arrisi.com

Mark Francisco
Comcast
mark_francisco@cable.comcast.com

Abstract

Adaptive delivery technologies which are in use to deliver services over unmanaged networks will allow for greater efficiency and capability of managed networks. A unified video processing workflow that can scale to large scale video delivery is described. Using this workflow from content ingest to distribution will provide a higher quality of experience over a wide range of devices. The concept of active management of adaptive delivery is proposed and described as a method of maintaining network capacity and incremental offering of new services and formats. A method of validating adaptive profiles is described and results presented comparing a number of different adaptive technology implementations.

OVERVIEW OF ADAPTIVE VIDEO DISTRIBUTION ARCHITECTURE

In this section, we present an overview of the architecture and basic video processing infrastructure used by an over-the-top video provider using Adaptive Bit Rate (ABR) video streaming. Delivery of video over an unmanaged network, like the public Internet, is based on best-effort delivery. Therefore the quality of user experience is impacted by various factors including the latency and congestion in the network, and uncontrolled variations of these network parameters over time. In spite of these limitations on-line video distribution has become prevalent recently and the quality of the video being distributed is ratcheting up slowly especially because of the increase in last-mile access bandwidth. Most on-line video distributors leverage some flavor of ABR streaming technology to adapt video delivery to vagaries of varying network conditions.

Most commercial over the top video delivery services (Netflix, Hulu, Amazon Prime, etc.) available today deliver on-demand content such as movies and TV shows, and not 24x7 live content. Exceptions to this are popular sporting events like the Olympics that are delivered over the Internet directly by the programmers. On the other hand, MVPDs (multichannel video programming distributor) have started to deliver live content using ABR (adaptive bitrate) streaming techniques to secondary viewing devices both inside and outside the home details of which will be presented in the subsequent sections.

Figure 1 shows a general content processing and delivery architecture employed by ABR streaming services. The first three elements, namely the Transcode, Package and Encrypt/DRM perform functions are related to content processing. The transcoder generates multiple *profiles* of video at different rates and resolutions, using H.264 for video and AAC (Advanced Audio Codec) for audio. Migration to more efficient codecs such as HEVC is expected to happen in the near future, but almost all streaming services currently use H.264 as the video format. Subsequent to transcoding, the packager encapsulates the video and audio elementary streams into one of many different container formats. The container formats are MPEG-2 Transport Streams (TS) for HLS, and variants of fragmented MPEG-4 file for HSS and HDS formats. Recently announced ISO standard DASH container format can provide a non-proprietary alternative to the above formats depending on how quickly DASH sees adoption commercially. As ABR streaming has moved towards delivering HD and premium content, content encryption and digital rights management (DRM) have become integral part of a successful delivery service as shown in the figure.

Following the content processing modules are the Origin and Edge cache servers that are responsible for eventual delivery of the processed content to the end clients. Large scale content delivery is performed using a multi-level caching architecture containing one or more high storage capacity Origin servers and a large number of edge cache servers with high streaming capacities. Over the top services employ third-party CDN (Content Distribution Networks) to deliver video. The edge cache servers in the CDN help to move the content closer to the end subscriber as well as caching the most popular content for repeated consumption by other users in the same geographical area. Efficient delivery of content over the Internet with acceptable of quality of user experience continues to be a challenge as the display and processing capabilities of the devices consuming video continue to improve rapidly.

ADAPTIVE VIDEO DELIVERY IN A MANAGED NETWORK

One of the advantages of the ABR streaming discussed in the previous section is that it can be used deliver delay sensitive content such as video over the Internet using best-effort delivery. The flow-control mechanism in current ABR streaming is initiated and managed by client devices. Software in client devices dynamically request different profiles based on various factors, such as the estimate of available bandwidth, network latency, decoding buffer fullness or instantaneous CPU usage. This approach, although found to work reasonably well for unmanaged delivery, results in a non-optimal usage of network resources and end user quality of experience as described below.

Sub-optimal usage of network resources can be illustrated with a rather simple example. Assume that an ABR service has three profiles, a 5 Mbps profile for delivery to a STB-connected large screen TV, a 4 Mbps profile for delivery to PCs and a 3 Mbps profile for delivery to tablet devices. Individual clients are allowed to request any of the aforementioned profiles. Let us assume that there is an available bandwidth of 11 Mbps and two tablet devices initiate new sessions for the content. These two clients can progressively request higher bandwidth profiles and say they go up to the 4 Mbps profile consuming 8 Mbps of bandwidth. After this if a STB initiates a request, the available bandwidth is only 3 Mbps, thereby forcing the STB to deliver a less than optimal video to the large screen TV. Clearly individual clients in this scenario do not have knowledge of requests coming from other clients and therefore they make decisions in a rather greedy fashion. In an optimal delivery scenario, both tablet clients can deliver good quality user experience using the 3 Mbps profile leaving the 5 Mbps bandwidth to the STB.

In a managed network, optimal resource allocation can be accomplished by using a central controller that has information about all the adaptive clients that share the same pipe. Therefore the controller will have knowledge of all the requests coming from devices that are in the same node or service group. Based on the knowledge of the ABR profiles and the capabilities of the client devices, the central controller can make optimal decisions on bandwidth allocation to individual clients.

The central controller can improve the quality of delivery by leveraging several factors in addition to the information on the type of clients and their requests as mentioned above. The central controller can monitor the instantaneous network utilization and proactively respond to any anticipated network congestion. Impact of any network congestion can be distributed over all the clients in the service group rather than adversely affecting a handful of clients. It can also implement business rules on content bitrates, priority of content and levels of subscription of end-user to decide on what content profile to deliver to individual clients. Figure 2 below shows one of the representative implementation scenarios of the central controller. Note that the central controller can reside in several points in the network, either in the cloud, a CMTS, one or more edge cache servers, or in a gateway inside the home.

UNIFIED WORKFLOW

Modern video receivers can be serviced with a standard set of video formats using adaptive delivery methods. It is possible to support a wide range of device types and network conditions with a relatively small set of bitrate/resolution profiles. The maximum desired resolution for medium to large screens such as television, computer and tablet is nearly identical and all devices support the lower resolutions required by smaller screens such as smartphones and portable media players. To achieve greater network efficiency, it may be desired to limit smaller screens to a defined subset of profiles that exclude excessive bitrates. Using similar methods, it may be desired to limit large screens to a subset that exclude profiles providing too poor an experience. These limits can be implemented through player configuration or manifest conditioning.

Increasingly, a single version of an asset may be delivered as video on-demand over QAM and IP. In some cases, the same asset may also be downloaded using IP. Broadcast content may also deliver over both QAM and IP and may be captured as a file for delivery as streaming recordings or video on-demand. From a singular file or stream contribution content may be adapted to use cases including broadcast, IP-streaming, network recording, video on demand and download. A unified workflow allows a single asset to be converted for multiple use cases and device types. An important aspect is to limit the variations of content available for streaming or at rest to bitrate/resolution profiles in support of adaptive delivery methods. The primary function of unified workflow is to utilize a single contribution format and create a set of files or streams that satisfy the bitrate / resolution requirements of all device types and network capabilities.

A subsequent operation is performed as assets are delivered to package the content for the specifics of the device platform. Packaging operations may include file fragmentation, transport multiplexing, audio binding and identity binding of content security. Frequently referred to as just-in-time packaging, these operations are practical to implement in high stream capacity appliances that are a component of the content origin section of the CDN. An end to end illustration of unified workflow is included in Figure 3.

DYNAMIC SELECTION OF ABR PROFILES

As shown in Figure 1, one of the important design considerations in ABR delivery is the selection of video profiles to be generated by the transcoder. Choosing a large number of profiles with small increments in bitrates and resolutions will enable client devices to smoothly switch between different profiles. However having a large number of profiles increases the complexity and cost associated with generating these profiles, maintaining the assets for future play-out and their ingest into origin and edge cache servers.

Table 1 shows a representative set of profiles used in an ABR service. These profiles, with their corresponding resolutions and bitrates, have been chosen *apriory* to match the capabilities of various display devices that are expected to have access to the content.

Profile	**Resolution (W x H)**	**Video Bitrate in mbps**
1	1280 x 720	3.0
2	1280 x 720	2.0
3	960 x 540	1.5
4	864 x 486	1.25
5	640 x 360	0.75
6	416 x 240	0.5
7	320 x 170	0.35

Table 1. Representative Set of Profiles

One of the challenges with hand-picking the profiles is that the selected bitrates and resolutions may be either too aggressive or conservative due to variability of the video content. We used a proprietary video quality tool to study the level of variability that can be encountered in practice. This tool estimates the perceived quality of ABR video segments and it was run on several different live video programs. Each of the video programs was coded at the profiles shown in Table 1 with segment duration of 6 seconds. The video quality tool that provides a no-reference scores for each segment in a scale of 1 to 100. A quality score of 75 or above is found to be visually acceptable based on our experiments. Figure 4 shows the plot of measured video quality of two different video content materials, "Home & Garden" and "CNBC" using Profile 1 encoding parameters. The "CNBC" video sequence exhibits a very high quality at Profile 1 encoding parameters, whereas the "Home & Garden" sequence exhibits larger variations and lower quality. This shows the difficulty associated with static selection of profiles.

We also investigated how the video quality scores change for a given live video program with different profiles used in Table 1. In some of the sequences, the video quality of different profiles is fairly evenly spaced apart, whereas for many sequences several profiles have similar video quality. Figure 5 shows the video quality score for the first three profiles of the CNBC sequence. As expected Profile 1 has a very high video score. Profiles 2 and 3 have a fairly similar video quality score indicating that Profile 3 could have been coded at a lower rate.

As the ABR video delivery technology matures, we believe that the profiles can be selected in a dynamic fashion in real time. A transcoder/encoder device that uses a two-pass video processing architecture can derive estimates on expected video quality for different profiles and then dynamically choose the combinations of profiles in a content adaptive fashion. In an alternate architecture, a transcoder can generate bitstreams corresponding to a large number of profiles. This can be followed by a downstream analysis device that inspects and compares the bitstreams corresponding to the individual profiles and then dynamically chooses a subset of profiles to be used. This can be accomplished in real-time with a small amount of processing delay.

<u>Profile selection - Adaptive performance comparisons across adaptive technologies</u>

The selection of encoding profiles is a key determiner of quality of experience for the customer and network capacity for the operator. Too few or improperly spaced adaptive variants may cause discontinuities in playback and marked video quality changes. Too many profiles increases the network storage requirements for a single asset and may cause unnecessary chattering between similar profiles as the adaptive player makes decisions based on available profiles and buffer condition. Encoding profiles are selected through a variety of rules of thumb and best practices. These include offering sufficient bits/pixel (typically 2-4), resolutions that are mod16 for efficient macroblock encoding, square pixel aspect ratio to ensure content is not anamorphically scaled on a variety of devices, and bitrate ratios between adaptive variants that are typically 1.5 – 2.

These rules of thumb can be tested by subjecting a player to dynamic network impairments while conducting subjective viewing of video quality and objective analysis of variant selection. Figure 6 illustrates an adaptive test system that was used to compare behavior of various adaptive technologies to a similar set of encoding profiles and network variations.

A number of adaptive delivery technologies are available and differ in platform support, multiplexing and manifest syntax. The performance of adaptive delivery is not determined by these variations, but mostly by the implementation of adaptive heuristics in the player. Three different adaptive technologies were evaluated; HTTP Live Streaming (HLS), HTTP Dynamic Streaming (HDS), and Smooth Streaming (MS Smooth) Each player was subject to the same time-based network variation and was provided an identical set of content encoding profiles to select from. The content profiles used are shown in Table 2.

Profile	**Resolution (W x H)**	**Video Bitrate in mbps**
1	1280 x 720	6.1
2	1280 x 720	3.3
3	768 x 432	1.5
4	640 x 360	1.1
5	512 x 288	0.85

Table 2. Profiles used for adaptive performance evaluation

The results in Figures 7 through 9 show how the unique decision methods implemented in each adaptive technology player affect the ability to react to variations in network condition. During the testing, no video discontinuities due to buffering or decoder starving were noted. In general, a player which maintains the highest profile selection and switches less frequently is preferred to one that selects lower resolution profiles and/or switches profiles frequently. The results suggest under slowly changing network bandwidth conditions, HDS and HLS maintain higher profiles than MS Smooth while HLS switches profiles less frequently. Each player was subjected to a second scenario of quickly changing network conditions with results shown in Figures 10 through 12. The results provide similar conclusions that HDS and HLS maintain higher profiles than MS Smooth while HLS switches less frequently. Video artifacts were noted occasionally on all three players during the intervals of greatest network congestion, as noted on the charts. This suggests a lower bitrate profile might be desired to improve performance under quickly deteriorating bandwidth situations.

<u>Adaptive performance comparisons across devices</u>

Adaptive delivery is customarily applied to address variations in network conditions to maintain a best possible experience. Adaptive delivery can also address variations in device performance. Figure 13 illustrates the response of devices with a variety of performance characteristics to an adaptive stream. The test environment shown in figure 6 was used to subject a number of devices to the same content profiles, adaptive technology and time-varying network variation. Each device was monitored to determine which content profile was selected during these variations to evaluate device specific behavior. The results suggest that most devices are more than suited for a wide range of content profiles and network variations. The least capable device in the test, the iPod Touch, was capable of selecting profiles normally targeted for large screen viewing, although the device switched away from the profile with less network congestion than higher powered devices.

THE ROLE OF ADAPTIVE DELIVERY WHEN OVERLAYING NEW FORMATS

A concern with unifying the formats among devices is the ability to introduce new profiles. Example applications may include advancing resolutions such as ultra high-definition (4K or 2160p), unique audio formats such as 7.1 or new video compression technologies such as HEVC. New video formats can be introduced incrementally by adding a variant to the adaptive encoding profile set. The set of profiles available to a device may be tailored to device capabilities through manifest conditioning or video player settings. A just-in-time packager is suited to introduce new audio formats by combining the appropriate audio and video formats for the device at time of delivery. This method can also be applied to introduce additional audio tracks such as alternate language or visually impaired commentary without burdening all streams with additive audio bandwidth.

SUMMARY

Traditional delivery of video over QAM infrastructure has been optimized over the last two decades. Adaptive IP video delivery, especially over a managed network, is still in its infancy. There are several challenges that exist in this new delivery paradigm and an array of new technologies that can be leveraged to improve IP video delivery. Using a unified work-flow for content ingestion will ensure that the content is generated at the highest possible quality and at resolutions that are consistent with various client devices. Experimental results presented in this paper show that clients respond somewhat differently to varying network conditions. Use of active management can ensure that the quality of experience can be judiciously and fairly maintained for all active users irrespective of client heuristics.

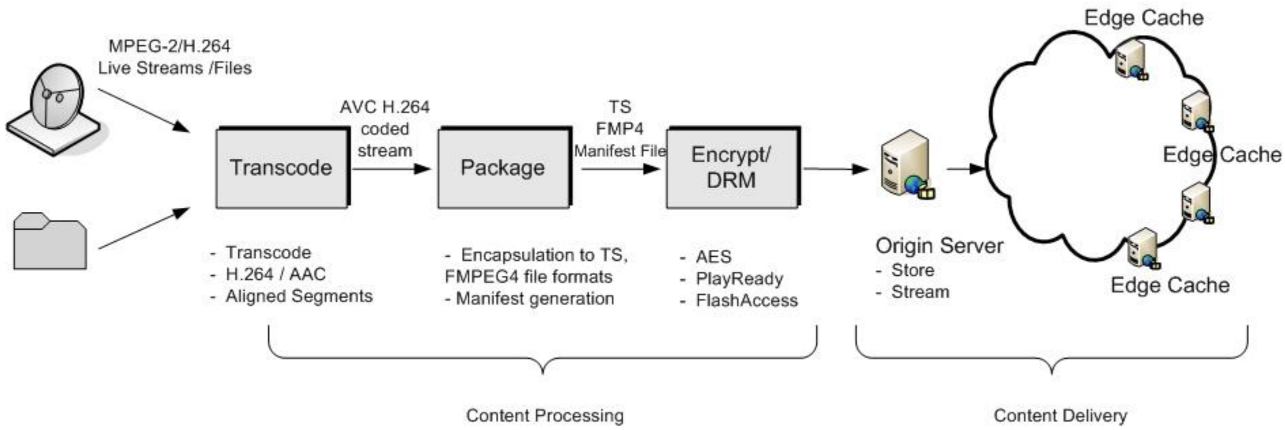

Figure 1. General Content Processing and Delivery Architecture

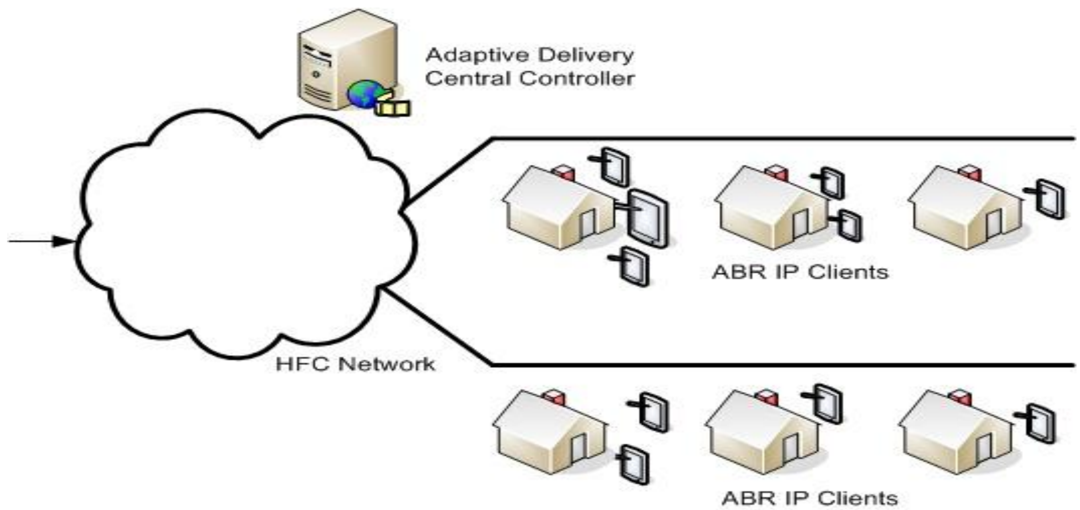

Figure 2. Central Control Representative Implementation

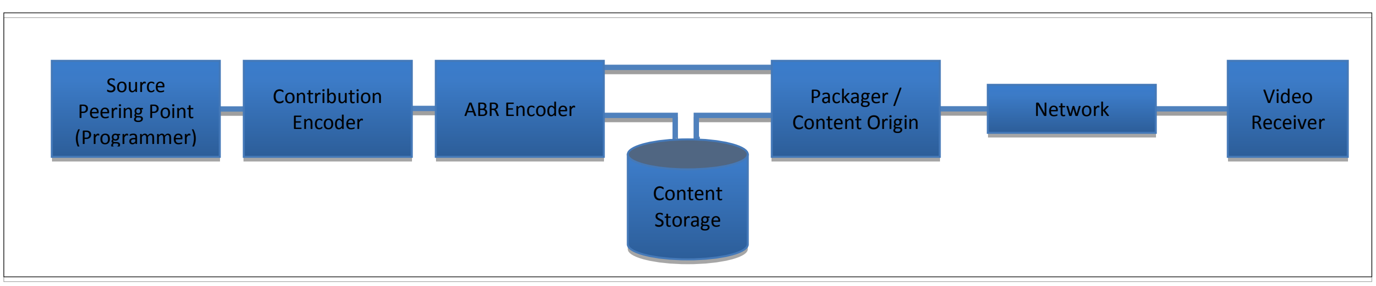

Figure 3. Unified Workflow

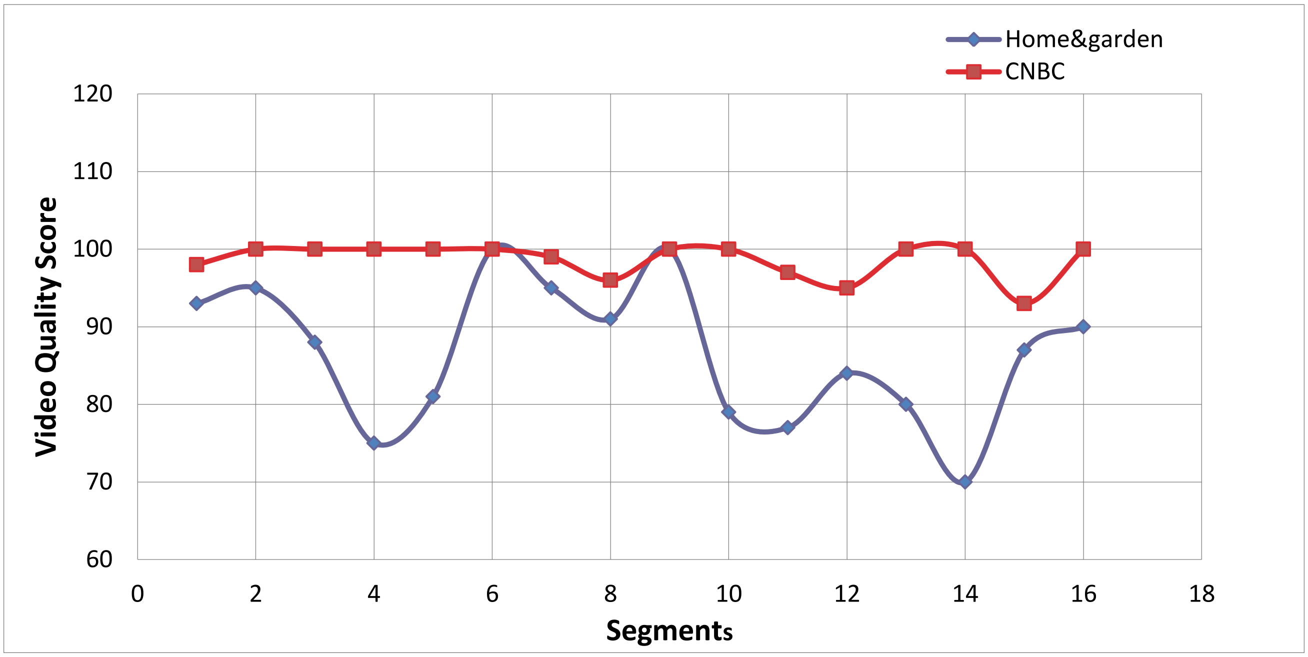

Figure 4. Video Quality with Table 1 Encoding Profiles

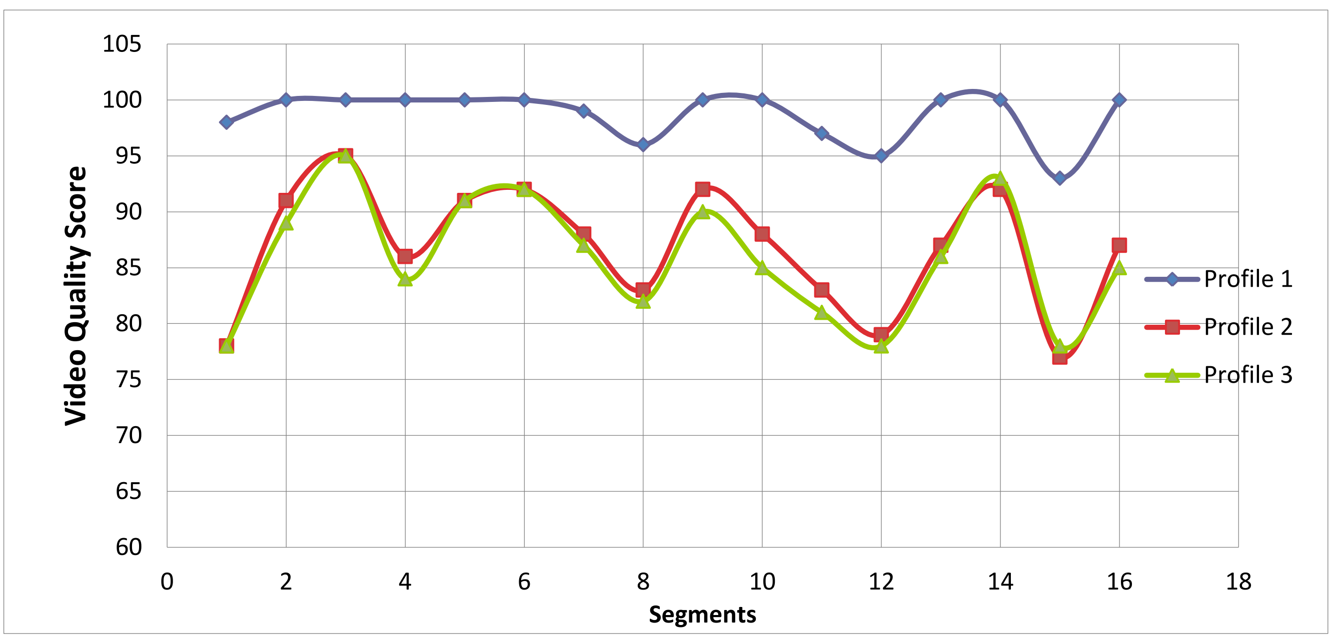

Figure 5. Video Quality Score for First Three Profiles

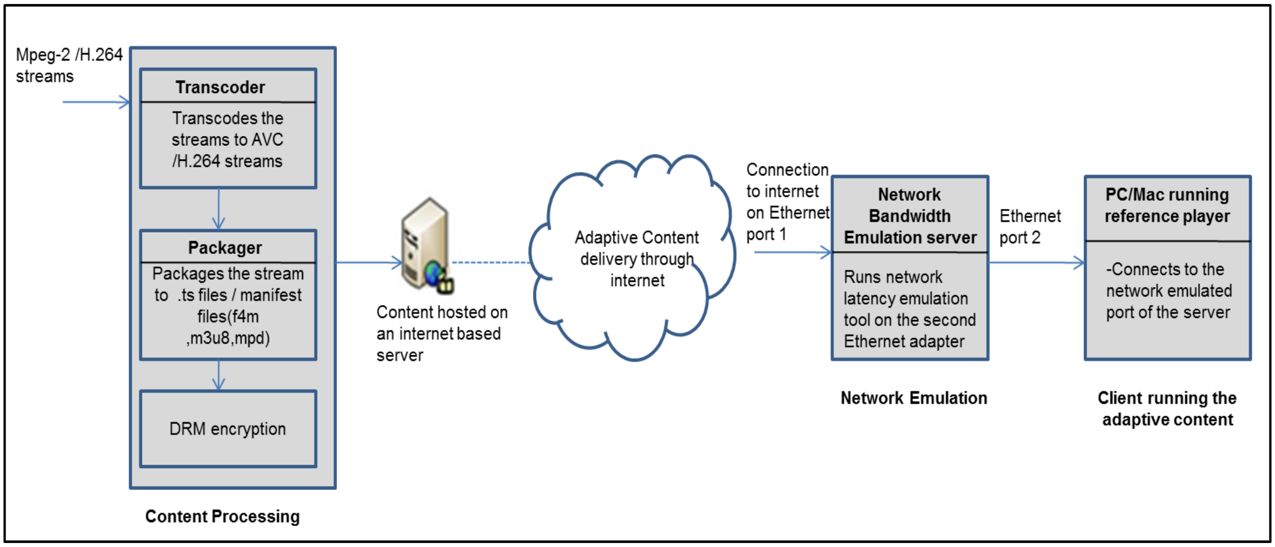

Figure 6. Adaptive Performance Test Environment

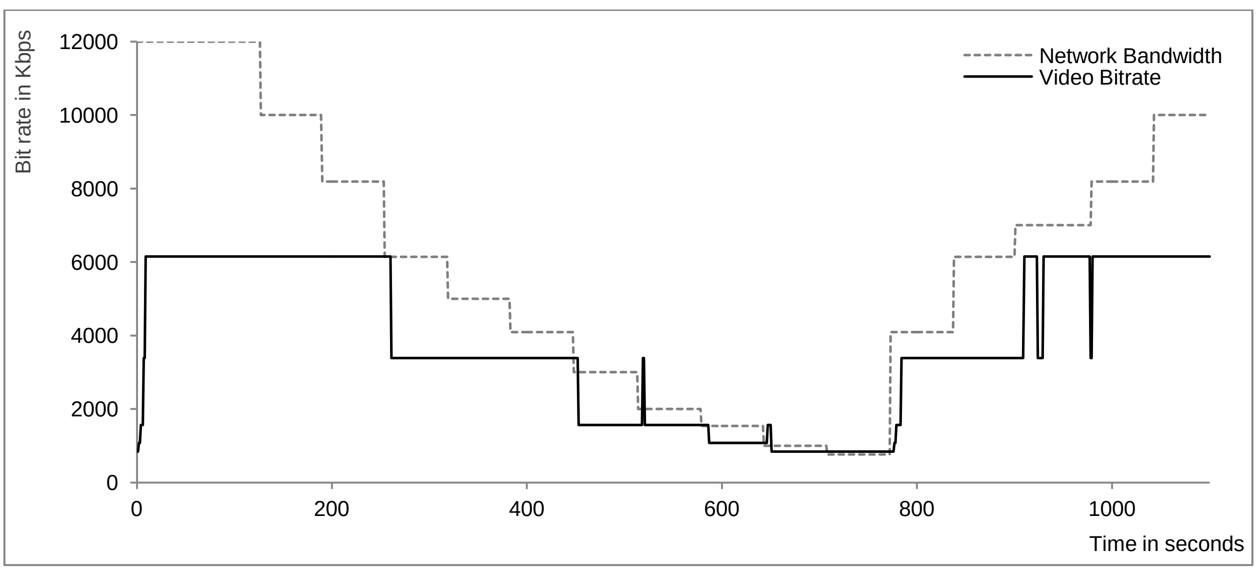

Figure 7. Adaptive Content Performance to Slow Variation - HDS

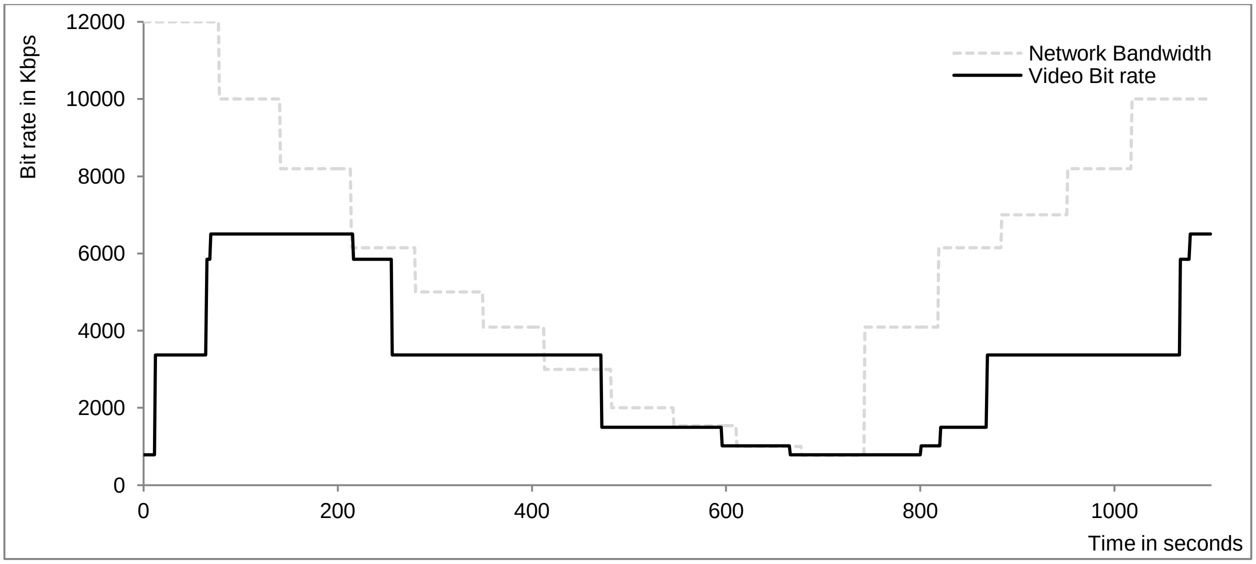

Figure 8. Adaptive Content Performance to Slow Variation - HLS

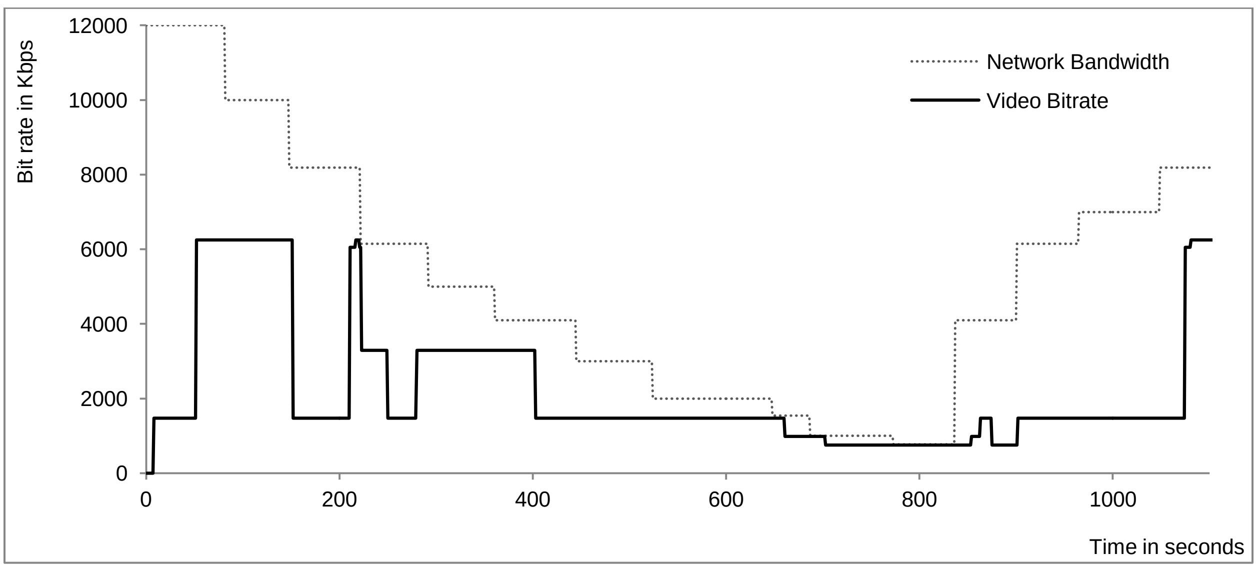

Figure 9. Adaptive Content Performance to Slow Variation -MS Smooth

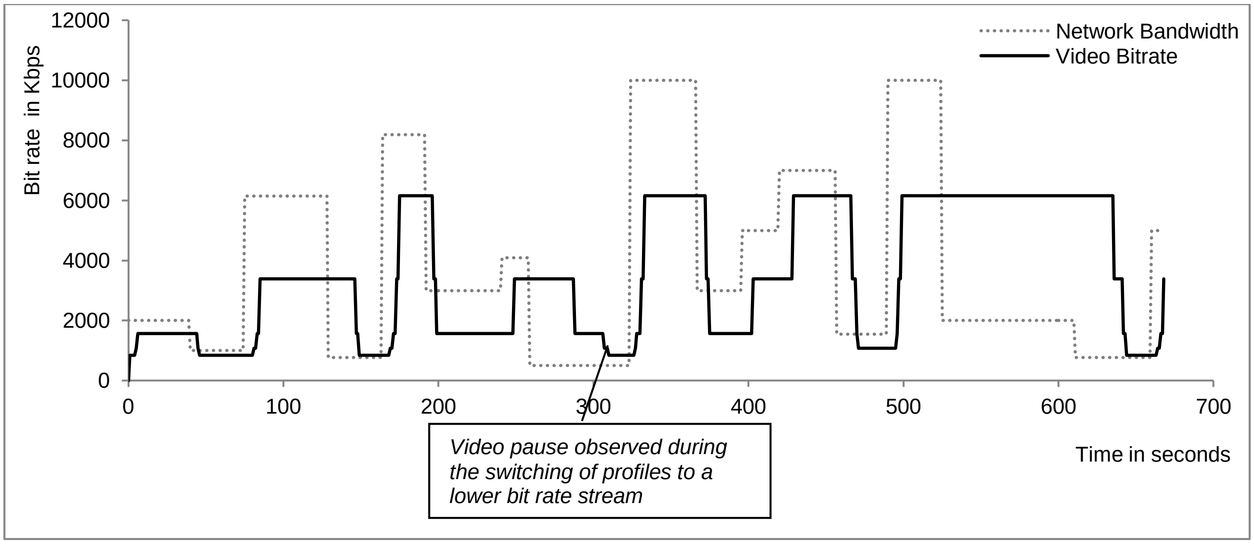

Figure 10. Adaptive response to rapid changes in network bandwidth - HDS

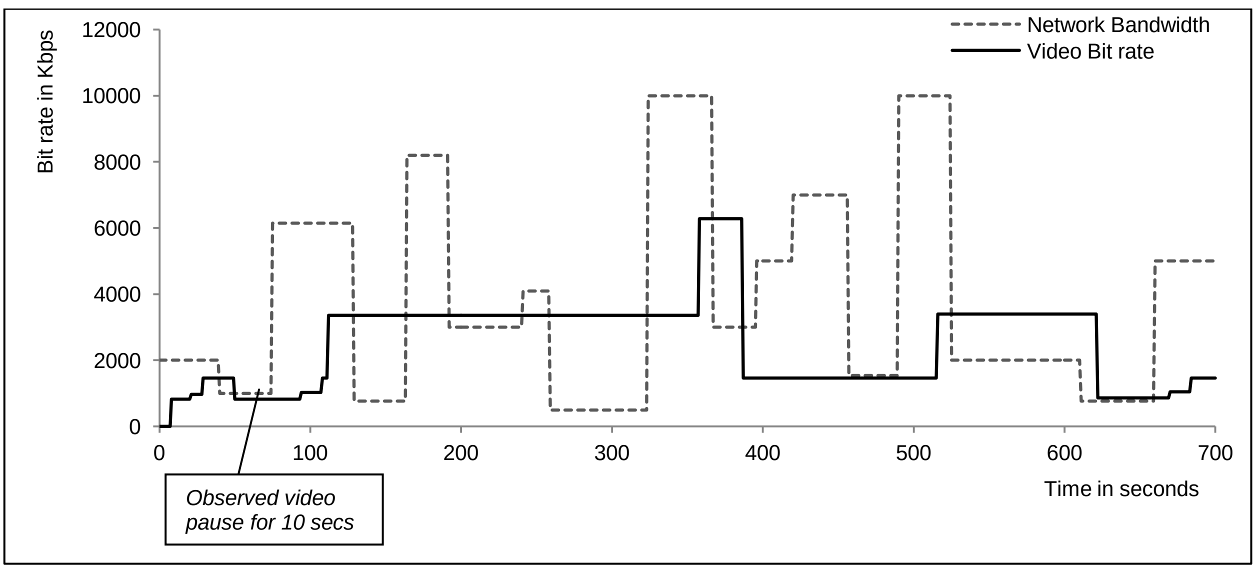

Figure 11. Adaptive response to rapid changes in network bandwidth – HLS

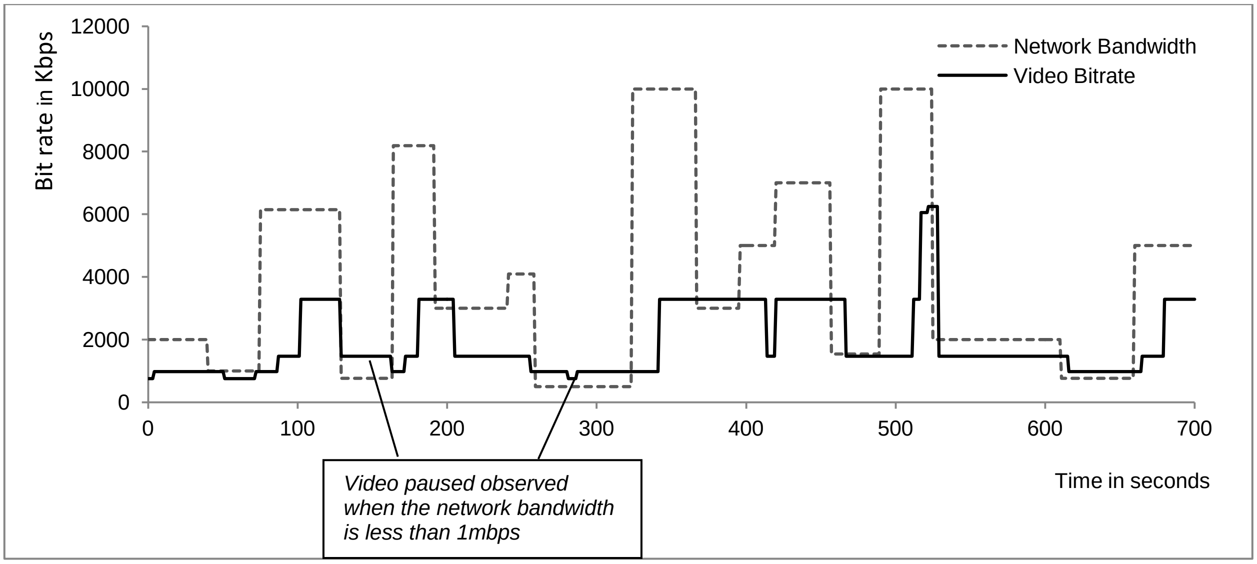

Figure 12. Adaptive response to rapid changes in network bandwidth – MS Smooth

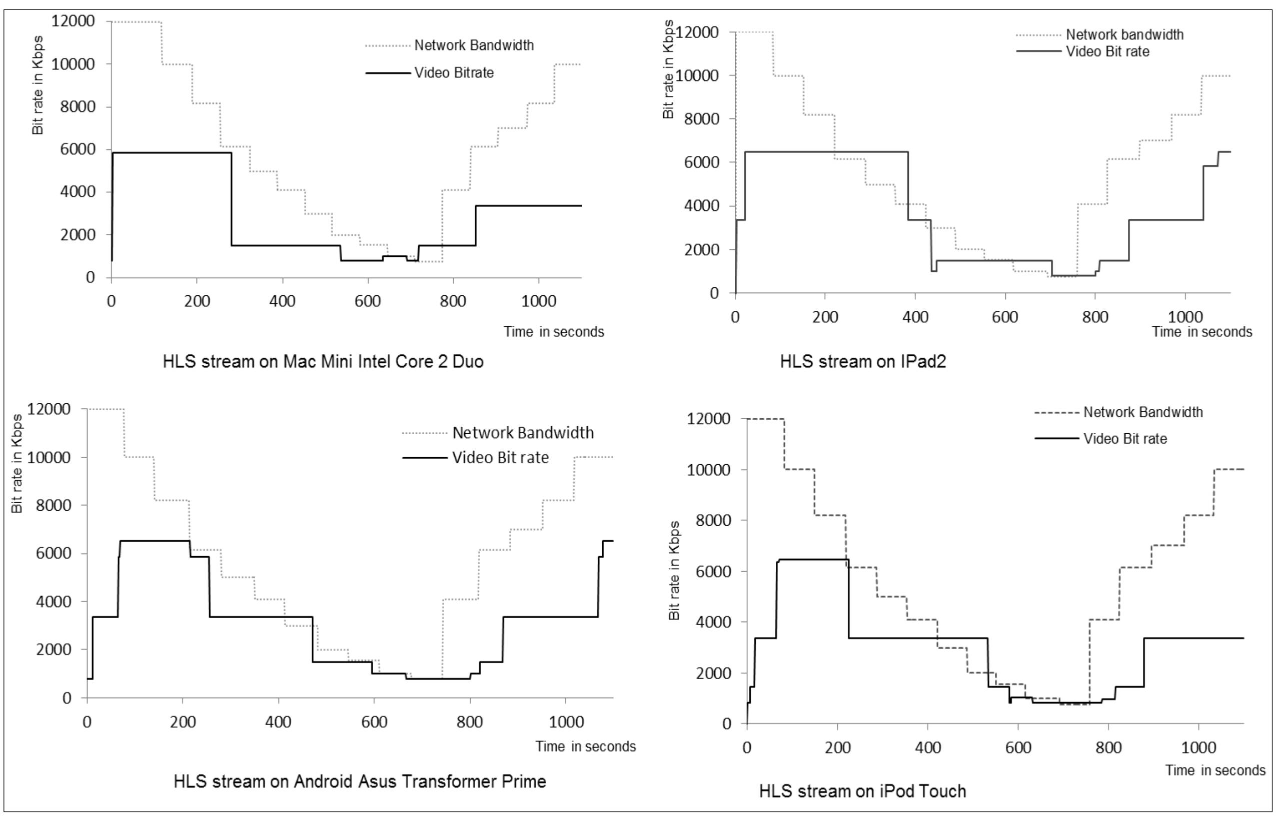

Figure 13. Adaptation to Device Capability

IMPROVING HOME NETWORKING SATISFACTION WITH A UNIFIED HOME GATEWAY

Carol Ansley

ARRIS

Abstract

Home users are becoming more and more dependent on their home IP networks to provide services and entertainment. As the industry moves toward a unified IP infrastructure for both HSD and video, the home network will be stressed by the increasing demands of video streaming to an ever-increasing universe of devices. The home network also already supports connections for computing devices such as PCs or laptops, and service devices such as scanners or printers. Increasingly the network also must support other devices such as tablets and smart phones in WiFi mode, with more devices showing up every week. Security and Smart Grid usage monitoring and control are also likely to become more and more common, as well as other services that are still on the drawing board.

This paper discusses the variety of networking technologies used to support the current home network, and proposes strategies to improve the customer's home networking experience with a three-pronged analysis.

First, what improvements can an in-home gateway (GW) provide by coordinating traffic across the wired and wireless networks? For example, if some devices can access more than one network, the GW could influence the device's choice by enacting QoS policies on a dual band device that segregate faster devices onto one band, and slower devices onto another band.

Second, can the user be encouraged to take actions that improve the home networking experience? For example, a GW could alert the user to an intermittent WiFi problem that has a time-based signature by monitoring wireless throughput over time.

Third, if the user calls a service provider with a networking complaint, can the GW or a network monitoring system supply a CSR with enough information to point out the most likely cause or causes of the user's problem? For example, a GW might flag to the user that one device is taking up 75% of a network's bandwidth because of poor choices in network setup.

CHALLENGES OF HOME NETWORK COMPLEXITY

Not so long ago, a home network might have been a small Ethernet bridge that connected together a computer, printer, and a cable modem. As wireless home networking became less expensive and easier for a consumer to install and configure, 802.11-based networks popped up. In the current home networking scene, the list of viable home networking technologies has gotten long. Some technologies, like MoCA, have been primarily used by service operators while others like Bluetooth or Zigbee have been driven by the consumer electronics industry.

The modern home is a fabric of networks with different capabilities that overlay each other. The following table compares some networks commonly found in a North American home.

Network	Wired/Wireless	Throuhput	Common Uses	Installer
Ethernet	*Cat5*	*Up to 1Gigabit*	*Data, Video*	*Consumer*
WiFi	*2.4GHz, 5GHz*	*Varies with distance, interference, channel width*	*Data*	*Mix*
HPNA	*Phone lines*	*?200Mb/s, but can vary*	*Data, Video*	*Service Provider*
MoCA	*Coax*	*Up to 400Mb/s*	*Video*	*Service Provider*
Bluetooth	*Short range, 2.4GHz*	*2Mb/s*	*Voice, Home Control*	*Consumer*
Zigbee	*2.4GHz*	*20-250kb/s*	*Home Control*	*Consumer*
DECT	*1.9GHz*	*32-500kb/s*	*Voice*	*Consumer*

Each of the networks above has a type of transaction for which it is best suited. The different networks also consume varying amounts of power in operation or when idle. As consumers become more conscious of energy usage, but still demand performance, the management of in-home networks needs to become more intelligent.

COORDINATION BETWEEN NETWORKS

When a Gateway is deployed in a household, it can act as the center of multiple layers of home networking. Most Gateway devices have a built-in router, often called an eRouter, to direct and police traffic among the various networks. First a review of the popular home networking technologies is appropriate.

Wireless Networking

The eRouter usually controls one or more wireless network Access Points (AP). Even though 802.11 WiFi standards have expanded into new frequency bands, consumers often do not know or understand that their WiFi devices may use different non-overlapping frequency bands. In North America, 2.4GHz and 5GHz are both allowed for unlicensed transmissions. The older

2.4GHz band is the most heavily used in most areas, but as new devices are purchased that support the 5GHz band, it is expected to gradually fill up as well.

Within the 802.11 standards are many tools that can be used to enforce Quality of Service (QoS) expectations. A service might be set up on one RF band or the other. The eRouter through the AP setup might have several Service Set Identifiers or SSIDs. An SSID can act like a sub-network within the larger WiFi channel. A device that wanted to stream video, for example, might be directed to use a specific SSID that had a higher priority within the eRouter than another SSID that was used for best effort data. Many APs can also monitor their local RF signals, and determine the best RF channels within each RF band. The “best channel” determination is not a static decision, but an ongoing process that must continue as long as there is activity on the network. By monitoring traffic patterns, monitoring RF activity, or by prior provisioning, a Gateway’s eRouter and AP can direct traffic across the WiFi networks to most efficiently use the available RF resources. A WiFi Alliance initiative called Wireless Multimedia Extentions (WME or WMM) can also be used to provide prioritized treatment to different types of traffic, such as video vs. best effort data.

Ethernet Networking

Most Gateway devices still have Ethernet ports. Older devices may still have 10/100 ports, but the most modern Gateways will have Gigabit Ethernet ports. The raw throughput and reliability of wired connections are still valued by many consumers. The eRouter can apply policies and QoS standards to the Ethernet ports, just as it can to the other networks. Ethernet also has the advantage of providing a consistently low power method of passing data, as compared to WiFi or MoCA.

HPNA/G.hn Networking

A less common home networking option, HPNA, makes use of the home telephone wiring for data transmission. The ITU G.hn specification also can use twisted pair wiring or powerline connections or coax networks,

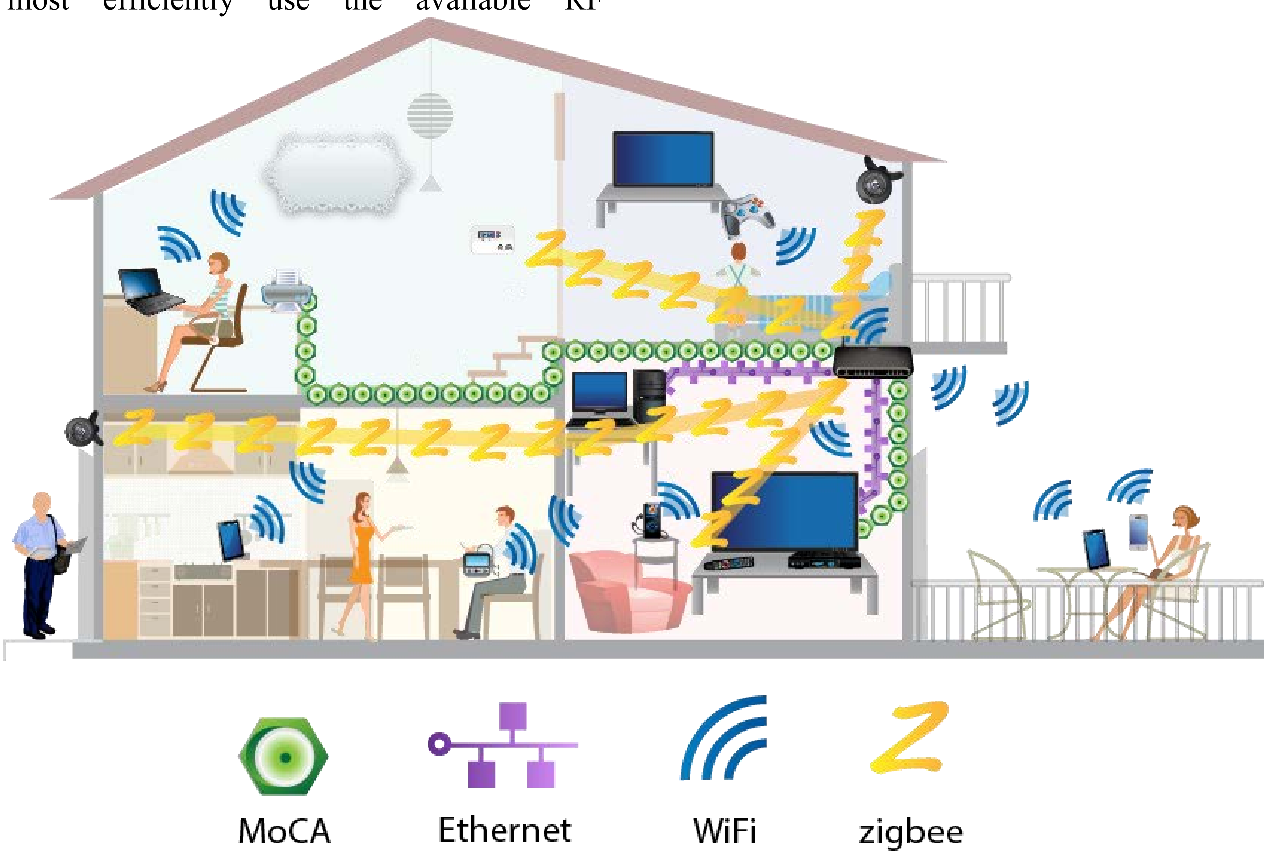

but it is not commonly found in CATV gateways.

MoCA Networking

The Multimedia Over Coax Alliance standardized a home networking technology that can utilize the home coax network without interfering with DOCSIS or analog video transmissions. MoCA home networking has primarily been used by the MSOs and Telcos for video distribution but as the price of consumer products comes down gradually, its use is broadening from just operator video to more general high speed data distribution within the house.

Similar to 802.11 WiFi, MoCA has more than one version available with different capabilities, using different amounts of RF bandwidth and providing different amounts of data throughput. A MoCA network operates within a certain frequency band; usually for MSOs, the D band between 1150 and 1650 MHz is used. Within that block of available bandwidth, the network can automatically settle on the best channel within that block of bandwidth or be directed to use a particular channel. Depending on the MoCA version, the channel bandwidth can change from 50MHz (MoCA1.1) to 100MHz (MoCA2.0) to 200MHz (MoCA2.0 with bonding).

MoCA has a QoS feature called pQoS, parameterized QoS. This feature allows filters to be set within the MoCA network to protect data flows. For example, an IP STB may require a dedicated pQoS flow for its video feed, but a MoCA/Ethernet bridge used for high speed data may just receive best effort treatment on the same network.

Bluetooth and Zigbee

Bluetooth and Zigbee services have not typically been considered part of the home data network, but as their popularity increases and the scope of services provided by operators increases, they will become more important. Bluetooth usage is usually very short range and relatively low bandwidth. Zigbee has been utilized for home control applications and for remote controls under the trade name RF4CE.

Coordination Across Networks

As consumer electronic devices become more sophisticated, a tendency has developed for many devices to have more than one choice of LAN interface. For instance, many printers and computers come standard with both wired and wireless interface choices. As the eRouter in a Gateway oversees network traffic within the home, it could act to improve the overall network performance by selectively directing devices to use one interface or another.

For example, a home network might have a laptop that can work in the 2.4GHz band or the 5GHz band, a MoCA to WiFi bridge that supports WiFi clients in the basement and 3 IP Settop Boxes (STBs) that can connect to the GW over MoCA or 5GHz wireless. The Gateway can balance the needs of the various clients as their activity occurs. During the day when the consumer is working in his basement office and the STBs are unused, the MoCA bridge could be allowed to have the majority of the MoCA wired bandwidth. The STBs could be directed to shut down their MoCA circuits to save power, and just listen over WiFi for any system updates. Later in the day as activity shifts in the household, the eRouter might monitor the non-video MoCA traffic. As STBs begin requesting video sessions, it would balance the amount of MoCA HSD traffic against the amount of 5GHz WiFi traffic and decide on a session by session basis if the STB should be served over WiFi or MoCA. There could be a fixed rule that the STBs get top priority on the MoCA link that would potentially lead to the HSD traffic over MOCA getting much poorer performance when video traffic was also competing for bandwidth. But, if the Gateway's eRouter can balance traffic across the networks by intelligently directing clients to the best interface, the same home networks can provide enhanced performance leading to a better customer experience.

In this example, the client devices had WiFi and MoCA networking available. Even if the second network is not capable of carrying much data, like Zigbee, it could still be used to communicate state information to remote clients. For example, a client with MoCA or WiFi and Zigbee could shut down the high bandwidth/high power MoCA or WiFi subsystem, and just watch over Zigbee for a wakeup indication. Zigbee was specifically deigned to be low power and low cost – it would make sense to utilize it more widely for control communications to save power.

USER INTERACTIONS

Consumers can also be given information by the Gateway to allow them to help improve their in-home network experience. Most shared media networks, such as MoCA or WiFi, can have their performance substantially degraded by the presence of clients conforming to older standards. The Gateway's eRouter can detect the presence of older technology on the home network, and make recommendations to the home user.

The recommendations might take the form of graphs showing how much bandwidth each connected device takes up typically on the network. The recommendations could also highlight user selections that may not be the best for performance. For example the user might have selected a particular WiFi frequency band based on the network characteristics when the system was being set up. The Gateway or eRouter might send an alert to the user indicating that the best

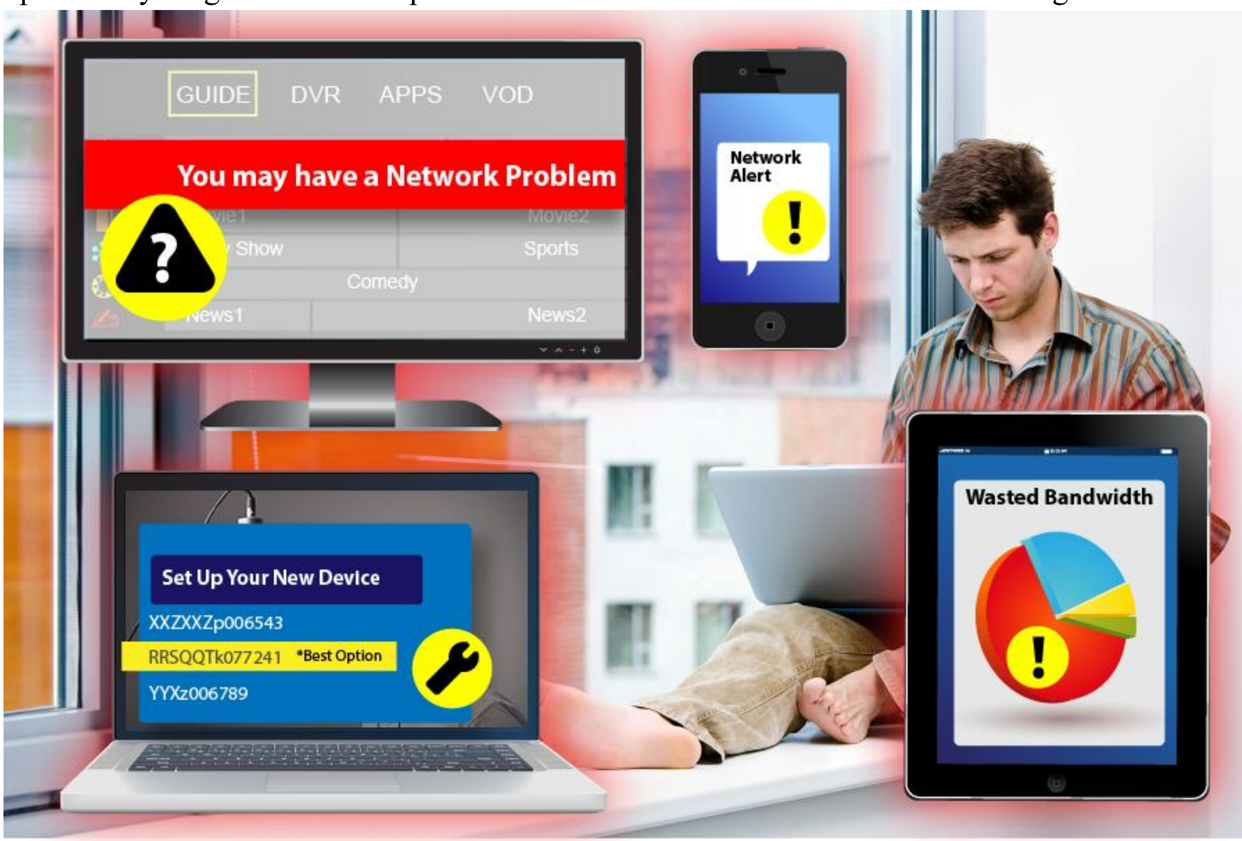

channel selection is now different and request that the user enable auto-selection or allow the Gateway to switch channels.

The user could be presented with recommendations by many means. The Gateway could email notifications to the consumer or post notifications to the user over social media. If a video service is being provided, the Gateway could pass user messages to the user interface application so that they could reach the user the next time they watch video programming.

The recommendations themselves might alert the user that one or more devices are slowing down the entire network with the consumer advised of some possible actions: take that device offline, change the offending device's settings to ones that will give the network better performance, or change some settings on the eRouter based on actual usage trends. Other actionable items could include reporting that some high-powered access points are showing up and causing interference. The consumer might be able to relate those itinerant APs to friends of his teenage children coming over.

NETWORK MAINTENANCE SUPPORT

The Gateway can provide information to the MSO's back office systems to enhance proactive network maintenance as well as interactions with the consumer. The eRouter within the Gateway can monitor conditions on various home networks and, through the Gateway's cable modem, the Gateway can also monitor its DOCSIS connection. The previous section dealt with using LAN information to advise the consumer of impending issues, or possible solutions to reduced network performance, but some problems may not be solvable by the consumer, or they may not understand enough technical details to troubleshoot the issue. This section discusses areas where the Gateway can assist in troubleshooting

Home LAN Troubleshooting

The Gateway can monitor, usually through the eRouter, all of the home networks. Depending upon the amount of storage available, it can record anomalous events such as a burst of strong interference on the 2.4GHz band, or loss of communications with a MoCA device. The event storage is most useful if it has time-stamps.

When a consumer calls a Customer Service Representative (CSR) with a problem, the Gateway can provide information to help track down the actual issue. For example, the consumer may be having difficulties with their wireless connection. Because wireless problems are often interference related, it is common for WiFi problems to come and go. If the Gateway has no historical information, then the consumer may experience frustration reporting a problem from the night before that is no longer evident in the morning. The CSR can only question the consumer about what might have happened and it is unlikely that the problem can be effectively resolved. If the Gateway records problematic events, then the CSR can see records of events that may be contributing to the poor performance experienced by the consumer.

The Gateway can also identify to the CSR items that may be affecting the user's home networking performance. Even if the Gateway also has provided feedback to the consumer directly, the consumer may still call a CSR for assistance understanding the problems and what actions can be taken to resolve them. If the CSR has access to the Gateway's recommendations and the data that it used to develop them, the CSR can explain the situation to the user more effectively, and hopefully ensure that the consumer has an excellent experience with swift resolution of any problems. For example, the Gateway might report that while the Gateway and other MSO clients are MoCA 2.0, the consumer may have purchased a MoCA1.x device that is lowering the performance of the network. Similarly, the wide variety of WiFi standards may result in the consumer not selecting the most effective settings in the WiFi network. If the user has allowed the Gateway to automatically configure itself, then the problem could have been resolved without the consumer's knowledge. If the options are not available to automatically configure the attached client devices, then the Gateway is dependent on the CSR explaining to the consumer how the settings are impacting network performance, and how to improve the network's performance.

IN SUMMARY

As MSOs deploy advanced Gateway devices that support multiple home networking technologies and have greater computing capacity than past generations of customer premises devices, the opportunity exists to utilize this technology to provide better home networking experiences to the consumer.

With multiple overlapping home network technologies in the Gateway and in some clients, the Gateway can take an active role in determining how to best utilize the networks for optimum performance and increased power savings. It can direct which network technologies are used and when by the attached devices to ensure the consumer's experience is optimized.

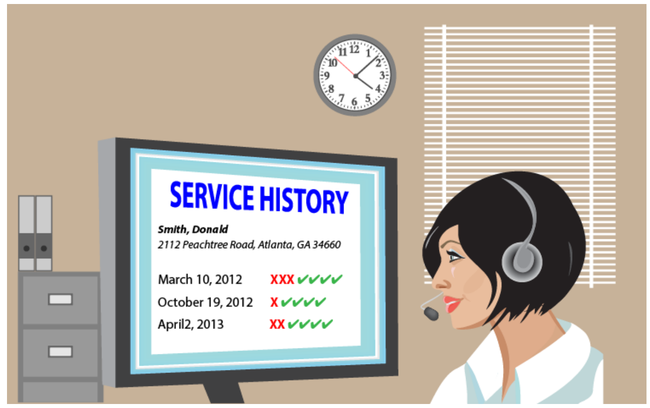

The unique position of the Gateway in the home can also allow it to provide feedback to help the consumer solve their own problems whenever possible. This capability has two benefits, the consumer gets better network performance without having to wait for an MSO technician to help them, and the MSO avoids spending valuable technician time on problems that the Gateway can diagnose and guide the consumer to solve independently.

Finally, for some problems, MSO involvement is unavoidable, but the Gateway can provide invaluable data to the operator that can pinpoint the most likely area at issue. The data may allow a CSR to guide the consumer through configuration options or other actions that will avoid a truck roll.

The new generation of advanced Gateway devices offers capabilities to provide the consumer with improved home networking performance and offers MSOs the potential for operational improvements and cost savings.

Is IPv6 over SP Wi-Fi a reality now?

Rajiv Asati, Distinguished Engineer (rajiva@cisco.com)
Ravi Shankar, Technical Leader (rshankar@cisco.com)
Cisco

Abstract

As Cable MSOs continue to exploit the fantastic opportunity with SP Wi-Fi and witness the exponential growth in their customer-base using the SP Wi-Fi services, they will soon have to decide whether to use public IPv4 addressing or private IPv4 addressing for the Wi-Fi Subscriber Equipments (UEs) used by their customer-base before the customers can avail of any free/paid services offered by the MSOs.

The reason is that most Cable MSOs have so far stuck with IPv4-only SP Wi-Fi designs/deployments, and most of them will run out of their public IPv4-address pools sooner or later (as soon as 2013-14 for some).

In this paper, we outline a strategy entailing IPv6 over SP Wi-Fi, discuss common challenges (that Cable MSOs face) with IPv6 in Wi-Fi and propose necessary solutions & technological evolutions for deploying IPv6 over SP Wi-Fi in Cable MSO networks.

TABLE OF CONTENTS

Table of Contents 1
Introduction 1
SP Wi-Fi architecture 2
SP Wi-Fi: IPv6 touchpoints 5
IPv6 Challenges & Solutions 6
- **IPv6 Address Assignment – SLAAC vs DHCPv6** 6
- **Multiple IPv6 Addresses with SLAAC** 6
- **Location Based Services** 7
- **Multiple IPv6 Addresses & IP Sessions** 8
- **Dual-Stack IP Sessions** 8
- **IPv6 Multicasting & IPv4-only Devices** 9
- **IPv6 Wi-Fi Client Mobility** 9
- **IPv6 Default Routing & Gateway Selection** 10
- **Transparent Auto-Logon (TAL)** 11
- **P2P Applications and First Hop Security (FHS)** 12
- **Lawful / Legal Intercept (LI)** 12
- **Prepaid Subscriptions** 13
- **Service Based Billing** 13

Evolutions 14
IPv6 Strategy / Recommendations 15
Summary 15
References 15

INTRODUCTION

Wi-Fi (an IEEE 802.11 family of standards [1]) is a pervasive & proven access technology that is now ubiquitously used indoors and outdoors. SP Wi-Fi, which primarily refers to an Wi-Fi system deployed and managed by a Service Provider (SP) for public access (aka community access) to SP network, is now commonly used by Cable Multi-System Operators (MSOs) to allow the customers/subscribers to benefit from the mobile consumption of Internet (and other) content. The MSOs are benefiting from SP

Wi-Fi to not only reduce the existing customer-attrition, but also attract new customers. MSOs can also offer the managed (and sometimes hosted) Wi-Fi services to other service providers (e.g. Mobile SPs).

As SP Wi-Fi continues to get widely deployed in MSO networks and get used more and more, MSOs must have to accommodate more and more devices/subscribers connecting to the networks, and use more and more IP addresses in addition to what MSOs already have/use for their triple-play services over wired access (e.g. Cable, Fiber).

> Imagine an SP Wi-Fi deployment having 10,000s of APs to serve 1,000,000s of subscriber devices. With 20% attachment rate, the MSO would need 210,000 IP addresses.

Noting the fact that almost all of the SP Wi-Fi deployments have been IPv4-only, and that many MSOs may not have enough IPv4 addresses left in very near future, it becomes obvious that operating SP Wi-Fi networks in an IPv4-only paradigm is quiet a significant business risk. While it is possible for MSOs to acquire additional IPv4 addresses from the RIR i.e. ARIN, they may not be able to do so once ARIN also runs out of its IPv4 addresses. Well, guess what, ARIN is already nearing the IPv4 address exhaustion [2].

This means something very obvious – Use IPv6. Of course, another obvious approach is to share each IPv4 address among number of customers.

MSOs must embrace/enable IPv6, no doubt about it. Of course, IPv6 enablement MUST NOT be done at the expense of IPv4. For all practical purposes, MSOs still need IPv4 (per customer or subscriber device) to ensure that IPv4-only devices and/or applications (aka apps) continue to work, despite having to dread the IPv4 address exhaustion. This means "Dual-stack IP addressing" (i.e. enable IPv6 while maintaining IPv4) of the customer/subscriber devices connecting to MSO Wi-Fi networks is MUST for now.

The key is to enable more and more subscribers/devices to use IPv6 for more and more content consumption, thereby limiting/reducing the need for IPv4, while ensuring customers' EXPERIENCE of content consumption.

SP WI-FI ARCHITECTURE

It is imperative to understand the SP Wi-Fi architecture as well as specific SP Wi-Fi components before dwelling into IPv6 specifics. This section provides a simplified overview of SP Wi-Fi architecture, which comprises of several building blocks:

1. Wi-Fi Access Points
2. Access Network
3. Metro/Aggregation Network
4. Wi-Fi Packet Core
5. Mobile Packet Core
6. Data Center

The Figure 1 illustrates such a simplified SP Wi-Fi architecture:

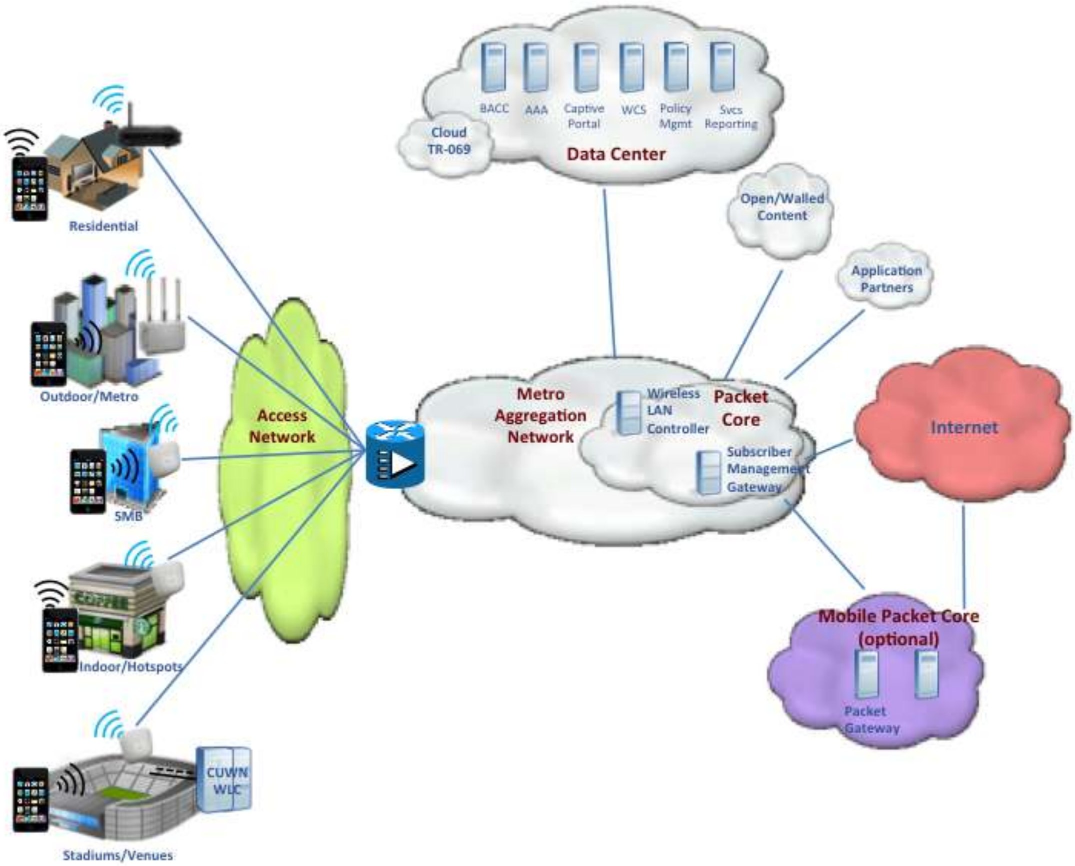

Figure 1 SP Wi-Fi Architecture – End-to-End (simplified)

The SP Wi-Fi architecture illustrated above contains the following building blocks:

1. Wi-Fi Access Points: The Wi-Fi Access Points may be either embedded with an ONU or cable modem (i.e. eDOCSIS device – also referred to as Cable Wi-Fi Gateway) or deployed separately from the CM or ONU.

2. Access Network: This is the DOCSIS based HFC network or xPON or Ethernet based Fiber network comprising CMTS or CCAP, Fiber Nodes, and CMs (or ONUs) providing network connectivity to/from the AP.

3. Metro/Aggregation Network: This is the network that CMTS uses to ultimately connect the subscribers to the internet or partner networks or the open/walled-garden content. There may also be a regional and/or backbone network (not shown in the figure) between the metro network and internet. Metro network is usually an IP or IP/MPLS network (or sometimes a layer2 Ethernet/bridged network).

4. Wi-Fi Packet Core: WPC typically comprises one or more Wireless LAN

Controllers (WLC) and Subscriber Management Gateways.

a. WLC – Responsible for control and management of Wi-Fi APs using CAPWAP protocol (IETF RFC 5415) and for mapping subscribers' traffic from Wi-Fi SSID to a virtual context e.g. VLAN by having tunnels to APs.

   Note that WLC is not expected to be present in the residential deployments of SP Wi-Fi.

b. Subscriber Management Gateway: An IP point of attachment that functions as a Policy Enforcement Point (PEP). Specifically, the gateway is responsible for maintaining subscriber awareness (in form of sessions), enforcing the per-subscriber policy e.g. QoS & bandwidth limits, and providing the usage/accounting, DPI, etc.

   The gateway is also referred to as Intelligent Services Gateway (ISG) or Session Manager (SM) or Gateway.

5. Data Center: The Service Network containing elements such as BAC, AAA, DNS, DHCP, Web-Portal, Policy Servers, OSS/BSS elements etc. providing network management and service management.

6. Mobile Packet Core: This is optional, but it is needed for ensuring inter-technology (4G/3G to Wi-Fi, say) or inter-domain mobility. This includes 3GPP specific elements such as PDN Gateway (PGW) etc. pertaining to cellular networks.

To understand the SP Wi-Fi components and their functions a bit better, we need to drill down in the above high-level end-to-end SP Wi-Fi architecture. Figure 2 focuses on the SP Wi-Fi components and their functions:

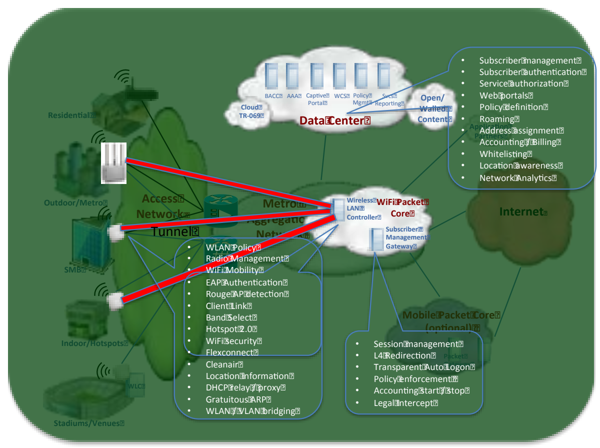

Figure 2 SP Wi-Fi Architecture - Functions

There are few points that are worth highlighting:

The WLC manages all the APs (including radio resources, SSIDs etc.) over the CAPWAP tunnels, and forwards/receives subscriber devices' IP traffic to/from Subscriber Management Gateway(s).

The WLC can either be L2 or L3 connected to the Subscriber Management Gateways. In either connectivity, all subscriber traffic must be forced through the Subscriber Management Gateway (so as to identify the subscriber, enforce policies, account for usage, perform Legal Interception etc.). To do so,

- In case of L2 connectivity, subscriber's default IP gateway is pointed to the Gateway's IP address
- In case of L3 connectivity, other IP routing mechanisms are used

The Subscriber Management Gateway terminates and manages subscriber sessions that are created dynamically when subscribers get online. Sessions are created as unauthenticated until subscriber credentials are verified and appropriate services are installed. Once authenticated, subscribers are allowed access into the network for the authorized content consumption.

SP WI-FI: IPV6 TOUCHPOINTS

IPv6 enablement in the SP Wi-Fi architecture can be divided in two categories – IPv6 enablement for Wi-Fi subscribers, and IPv6 enablement for Infrastructure.

The latter category is for enabling IPv6 in the network infrastructure independent of subscribers' usage of IPv6 (or IPv4). For example, AP and WLC communicating (CAPWAP) over IPv6.

The former category is for enabling IPv6 for the subscribers independent of the network infrastructure. This category requires a lot many more devices (besides the subscriber devices) in the Data Center to be enabled with IPv6. For example, AAA (or BSS) servers accommodating subscriber device' IPv6 address usage during authentication.

The figure below illustrates the IPv6 touch-points:

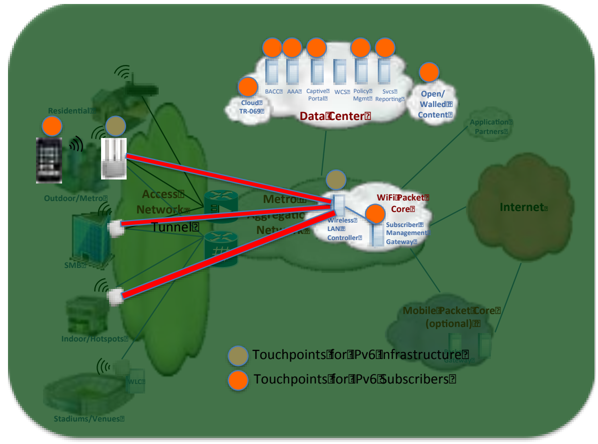

Figure 3 SP Wi-Fi Architecture: IPv6 Touchpoints

Unlike IPv4, IPv6 allows a device to have multiple IPv6 addresses. Hence, it is important to track all the IPv6 addresses per device and enable seamless mobility for them when the subscriber's device moves within the mobility domain.

> A mobility domain is defined as a contiguous area of Wi-Fi coverage in which a subscriber can move seamlessly and maintain IP session continuity.

The essential requirement for seamless mobility is that it should be transparent to the subscriber device (i.e. Wi-Fi Client) and it should allow the device to keep and use the original IP characteristics (IP address, default GW, DNS and DHCP services) after the move.

IPV6 CHALLENGES & SOLUTIONS

IPv6 enablement in the SP Wi-Fi architecture can lead to a number of challenges for any deployment. This section describes those challenges as well as current solutions.

IPv6 Address Assignment – SLAAC vs DHCPv6

In SP Wi-Fi, each subscriber device must be assigned at least one global IPv6 address in a dynamic manner. To do so, there are two options – SLAAC (Stateless Auto Address Configuration) and DHCPv6.

Challenge*: Use SLAAC or DHCPv6?*

SLAAC is stateless, hence, it scales better than DHCP(v6), which is stateful (address assignment). More importantly, SLAAC is ubiquitously supported on almost all subscriber devices, whereas DHCPv6 is not.

Despite its simplicity and ubiquitous support, SLAAC is not quite useful until unless one of the following two is also used for conveying DNS configuration information to the devices: (1) Router Advertisement (RA) option for DNS [RFC6106], (2) Stateless DHCPv6 [RFC3736]. At the time of writing this paper, the second option i.e. Stateless DHCPv6 has better support on various subscriber devices than that of the first option.

DHCPv6 support for address assignment on the subscriber devices is not as ubiquitous as that of SLAAC. For ex, many Android devices (e.g. Samsung Galaxy 7'' tablet with Android 4.1.0+) don't support DHCPv6 for address assignment. Nonetheless, DHCPv6 allows the addresses to be centrally managed for ongoing allocation to the subscribers.

So, which protocol to use in a deployment?

Solution*: Use SLAAC (along with Stateless DHCPv6) short-term/medium-term.*

It is worth pointing out that given the recommendation to enable dual-stacking on subscriber devices, they would be provided with the IPv4 DNS server addresses anyway (and DNS allows A and AAAA record look up over both IPv4 and IPv6), so IPv6 DNS server address is not mandatory to ensure IPv6 usage. This takes care of a small percentage of devices not supporting stateless DHCPv6.

The MSOs should resort to alternate methods such as (RADIUS) accounting start records to keep track of IPv6 addresses used by the subscriber devices for compliance & lawful intercept purposes.

Multiple IPv6 Addresses with SLAAC

When SLAAC is used for IPv6 address assignment, each subscriber device is provided (by the Gateway) an IPv6 prefix (of prefix-length = 64), which could be a dedicated to a device or shared among a number of devices.

Challenge*: Use Dedicated vs. Shared IPv6 Prefix?*

In case of a dedicated IPv6 prefix, each subscriber device gets a unique global IPv6 prefix and uses it to assign one or more global IPv6 addresses to itself.

In case of a shared IPv6 prefix, a number of subscriber devices get a common global IPv6 prefix and use it to assign one or more global IPv6 addresses.

The shared IPv6 prefix approach causes the Gateway to store each device's each IPv6 address in the routing/forwarding table, thereby increasing the size of the IPv6 routing table by 3-8 times (as a device may assign multiple IPv6 addresses). Moreover, as the device changes its IPv6 address (due to privacy extensions), the control plane traffic (Neighbor Discovery traffic corresponding to Duplicate Address Detection (DAD)) would unnecessarily consume the precious Wi-Fi bandwidth and cause the router to keep on updating the routing/forwarding table with more than one routes per device. This arguably is sub-optimal and unnecessary.

Solution: *Dedicated IPv6 prefix (of length = 64).*

A dedicated IPv6 prefix helps to keep a single routing/forwarding entry per device.

Also, it helps to identify an individual subscriber device's IPv6 traffic based on its prefix instead of any of its individual IPv6 addresses. This allows Gateway to use a single IP session per subscriber device and provide accounting simplicity.

Location Based Services

When SLAAC is used for IPv6 address assignment, each subscriber device is provided (by the Gateway) an IPv6 prefix (of prefix-length = 64), which could be a dedicated to a device or shared among a number of devices. In SP Wi-Fi, location based services rely on the hardcoded AP's location information to proximate the subscriber device's location

Challenge: *Location based services may fail with IPv6 when SLAAC is used for address assignment.*

In IPv4-only SP Wi-Fi, location based services can be offered by letting the web portal initiate a DHCPv4 lease query for the subscriber IPv4 address and receive the DHCP option 82 information (that would convey the corresponding AP location, as inserted by the WLC acting as a DHCP relay agent during the IPv4 address assignment to the subscriber device) to figure out the subscriber device's approximate location.

In IPv6 enabled SP Wi-Fi, location based services would not work, if DHCPv6 is not involved in address assignment. So, the portal cannot determine location based on source IPv6 address of the subscriber device.

Solution: *Determine Location based on VLAN and AP mapping.*

When subscriber devices are L2 connected to the 'Subscriber Management Gateway', all traffic from a set of APs can be mapped to specific VLANs. Location is then determined by virtue of incoming VLAN's by defining different captive portals per VLAN. This is illustrated in the figure below:

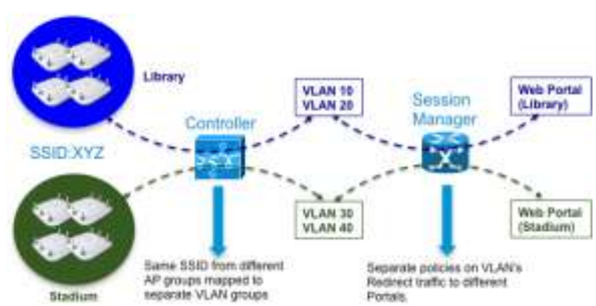

Figure 4 Location based Service - IPv6 Solution

The 'Subscriber Management Gateway' would need to have different redirection portals defined on incoming VLAN's and location relevance is based on redirection to the right instance of the portal.

Other options include out of band mechanisms for location relevance such as accounting start records sent from WLC to a database that have both the IPv6 address of the subscriber device and the associated AP name or location string.

Multiple IPv6 Addresses & IP Sessions

Unlike IPv4, IPv6 allows multiple IPv6 addresses per interface when an IPv6 prefix (shared or dedicated) is assigned to the subscriber device.

***Challenge**: Subscriber Management Gateway may create multiple IP sessions for the same subscriber device.*

IPv6 address that is used by the subscriber during the web authentication is associated with the 'session' on Subscriber Management Gateway and all traffic from the subscriber is identified based on this source IPv6 address. When the Wi-Fi Client has multiple IPv6 addresses and could potentially use different ones, it becomes necessary to associate all Wi-Fi Client IPv6 addresses with the existing session. Otherwise, the Gateway may create multiple sessions for the same device, impacting the overall session scale.

***Solution**: Use MAC address for creating a single IP session per subscriber device.*

A common unique identifier such as a device MAC address can be used to identify a session, but this requires that the Wi-Fi Client be L2 connected to the Subscriber Management Gateway, either directly or through a tunneling mechanism.

This is deemed advantageous for another benefit – having a single session and consolidated accounting records for both IPv4 and IPv6 traffic.

If the subscriber device can not be L2 connected, then an alternate mechanism where every subscriber device gets a dedicated IPv6 prefix could be utilized, so the traffic from the device can be identified based on the common prefix instead of individual IPv6 addresses. Needless to say that such a mechanism would ensure a single IP session for IPv6.

Dual-Stack IP Sessions

***Challenge**: Subscriber Management Gateway may not create a single IP session.*

Gateway may create separate IP sessions – one for IPv4 traffic and another for IPv6 traffic from/to the same subscriber device, depending on the Gateway implementations. This can impact the session scale.

***Solution**: Use MAC address for creating a single IP session per subscriber device*

A common unique identifier such as a device MAC address can be used to identify a session, but this requires that the Wi-Fi Client be L2 connected to the Subscriber Management Gateway, either directly or through a tunneling mechanism.

If the subscriber device can not be L2 connected, then this challenge can not be solved – two sessions must have to be created on the Gateway.

IPv6 Multicasting & IPv4-only Devices

IPv6 fundamentally relies on Multicasting such that the (Neighbor Discovery (including Router Discovery)) messages are sent to few or every subscriber device in the WLAN domain, thereby consuming the precious Wi-Fi radio resources.

Challenge*: Inefficient Wi-Fi radio utilization due to Multicasting.*

If the number of IPv6 enabled devices is less than that of IPv4-only devices in a given WLAN, then IPv6 multicasting to all nodes may sub-optimally consume Wi-Fi radio resources.

In IPv4 over SP Wi-Fi, the WLC usually acts as the DHCP relay and as the proxy for the subscriber devices' ARP requests. ARP requests for Wi-Fi Client MAC addresses are never sent across the backhaul network. Similar optimization is not available for IPv6.

Solution*: Use IPv6 unicasting for Router Advertisements and let WLC respond to NS messages using local ND cache.*

If all subscriber devices are not dual-stacked (e.g. IPv6 enabled), then it would be better to deliver IPv6 Router Advertisement messages (from Gateways) as IPv6 unicast messages only to those devices that are dual-stacked, so as to save the Wi-Fi radio resources as well as the backhaul infrastructure resources.

To do so, the APs and WLCs will have to be aware of the subscriber device type (IPv6 enabled or not) and unicast the RA messages to the IPv6 enabled devices. Additionally, just like with IPv4, the WLCs would have to be IPv6 aware and respond to NS (Neighbor Solicitation) messages coming from the devices using its local ND cache. NS messages for IPv6 addresses of devices are used by the Gateway as a keepalive mechanism to detect when the device has disconnected from the network. Subscriber management gateway uses NS messages for IPv6 addresses of clients as a keepalive mechanism. These messages are multicast by the subscriber management gateway, and must be responded to by the Wireless LAN controller instead of being sent downstream across the backhaul network and then across the radio network by the AP.

This yields better performance overall.

IPv6 Wi-Fi Client Mobility

It is commonly expected that some of the subscribers would move while connected to the Wi-Fi network and it would be desirable to not let their IP communication get interrupted due to the move.

Challenge*: IPv6 session persistency during the move within a mobility domain.*

Unlike IPv4 SP Wi-Fi network in which Discover Network Attachment (DNAv4) protocol is used by the Wi-Fi Clients to check its connectivity to the same L2 network during the move and ensure continuing with its existing IPv4 address, IPv6 has no such provisions.

With SLAAC based IPv6 address assignment, the Wi-Fi Client must have to receive the same prefix being advertised (in RA message) in order to maintain its IPv6 address. If the Wi-Fi Client sees a different prefix, then it will assign another new IPv6 address and deprecate the previous IPv6 address, eventually impacting the ongoing TCP/UDP sessions (due to upstream and/or downstream traffic not getting correctly forwarded by the Gateway to the Wi-Fi Client).

Solution: *Deliver the original RA (and IPv6 prefix) to the Wi-Fi Client after the move.*

When a Wi-Fi Client moves from one AP to the next, it will receive the same RA, if it is on the same (WLC-Gateway) VLAN after the move as before the move. However, after the move, if the Wi-Fi Client is attached to an AP on a different (WLC-Gateway) VLAN, then the Wi-Fi Client move will have to be tracked and the RA from the original VLAN will still have to be sent to the Wi-Fi Client so it can continue to use its existing address and via the previous default router, where the IP session existed.

This means that the traffic to and from the Wi-Fi Client, including ND messages, will have to be tunneled between the old and new WLCs to keep the Wi-Fi Client and the Gateway transparent to the move. This is possible because the WLCs in the same mobility domain exchange mobility messages to keep track of Wi-Fi Client movement. In case of L3 roams, the Wi-Fi Client remains anchored at its home WLC by tunneling all traffic to and from the Wi-Fi Client between the foreign and home WLC. The figure below shows IPv6 Wi-Fi Client mobility:

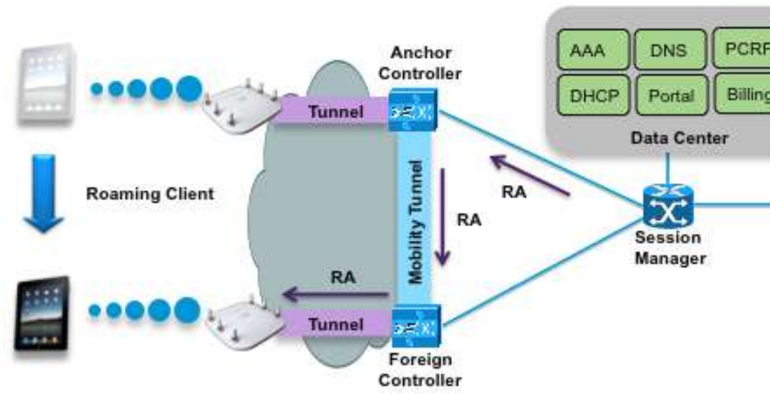

Figure 5 IPv6 SP Wi-Fi Client Mobility

The RA (Router Advertisement) is tunneled through the mobility tunnel between Anchor and Foreign WLC and unicasted to the Wi-Fi Client. Similarly, DHCPv6 and NDP messages are also processed at the anchor WLC through the dynamically established mobility tunnel.

IPv6 Default Routing & Gateway Selection

IPv6 routers (e.g. Subscriber Management Gateways) send ICMPv6 RA (Router Advertisement) messages to “all nodes” multicast address at a regular interval so as to inform all hosts of the router presence and characteristics, and to provide hosts with parameters they need to function properly on the network.

A host that wants to find out immediately what routers are present may send a RS (Router Solicitation), which will prompt listening routers to send out RA’s. RS is sent to the “all routers” multicast address. RA in response to a RS goes unicast back to the device that sent the RS, and is called a unicast RA.

__Challenge__: High-Availability and Load-sharing of Subscriber Management Gateways.

In IPv6 SP Wi-Fi deployment, the Wi-Fi Clients typically choose their default gateways based on router preference included in the Router Advertisements. When there are multiple RAs of equal priority, Wi-Fi Clients usually choose the router that sent the very first RA as their default gateway.

In SP Wi-Fi deployments, it is preferable to load balance Wi-Fi Clients across multiple Subscriber Management Gateways. Wi-Fi Clients sessions are established on the Subscriber Management Gateway based on their selection of the default gateway. In the case of IPv4, default gateway selection is enforced by DHCP configuration. With IPv6, the Wi-Fi Clients select their own default gateway and all subscriber sessions could end up on the same Subscriber Management Gateway.

__Solution__: Use VRRP with or without IPv6 default Router Priorities.

Set RA priorities (High, medium, low) on the Subscriber Management Gateways in order to influence the selection of default gateways. Wi-Fi Client traffic for the same SSID is load balanced on a per subscriber basis across a set of VLAN's. One Subscriber Management Gateway will have a "high" priority set for one or more VLAN's in the set and have a lower priority on the remaining VLANs. IPv6 traffic can be load balanced across a redundant set (N:1 or 1:1) of Subscriber Management Gateways.

The figure below shows an overview of load balancing and redundancy for IPv6 Wi-Fi Clients:

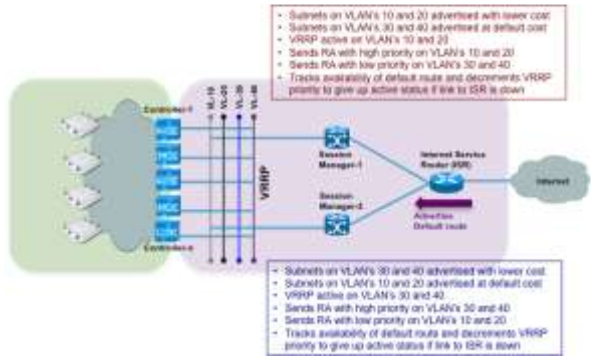

Figure 6 VRRP and Load-sharing

The same SSID is load balanced across a group of VLAN's on a per subscriber basis by the set of WLC's. Depending on the VLAN's assigned to Wi-Fi Clients, they will choose the Subscriber Management Gateway with the higher priority RA as their default router. Downstream traffic to the Wi-Fi Client will routed by the ISR (Internet Services Router) through the correct Subscriber Management Gateway based on route preferences at the ISR. When a subscriber-facing interface goes down, VRRP will ensure that Wi-Fi Client traffic will be processed by the redundant Subscriber Management Gateway and traffic symmetry (through the same Subscriber Management Gateway) will be maintained after route convergence at the ISR. Downtime is limited to VRRP and route convergence.

__Transparent Auto-Logon (TAL)__

Auto Logon (TAL) is used to enhance the subscriber experience by recognizing the Wi-Fi Client and bypassing the authentication requirement for a previously authenticated subscriber using the same device for a fixed time period, say 48 hours. As long as the subscriber accesses the network at least once during the set period, he/she will not have to authenticate on the network again since their identity is maintained in the network.

TAL is done by caching a unique ID – typically the MAC address of the Wi-Fi Client and then authorizing subsequent access against this cache based on the source MAC address of the incoming packet.

Challenge*: Multi-hop L3 connectivity between subscriber device and Gateway.*

When subscribers are L2 connected, the Subscriber Management Gateway has access to the Wi-Fi Client's MAC address and can attempt an authorization. If the Wi-Fi Client is L3 connected and using an IPv4 address, the Subscriber Management Gateway uses a DHCP lease query to obtain the MAC address of the Wi-Fi Client. When, the Wi-Fi Client is L3 connected and has an IPv6 address (Stateless or SLAAC), the Gateway has no visibility into the Wi-Fi Client MAC address and cannot attempt an authorization

Solution*: Enforce L2 connectivity or DHCPv6.*

To ensure TAL, IPv6 Wi-Fi Clients will have to be L2 connected to the Subscriber Management Gateway, so that they can be authorized using their MAC addresses. Alternatively, stateful DHCPv6 will have to be enforced so the Subscriber Management Gateway can use a DHCPv6 lease query to determine Wi-Fi Client MAC address for authorization.

P2P Applications and First Hop Security (FHS)

Peer-to-Peer communications must be allowed in a public SP Wi-Fi network to support many P2P applications like Skype, instant messaging, chat etc.

Challenge*: P2P communications allowance may also allow Layer 2 connectivity among the Wi-Fi clients.*

If any WLC implementations require having all Wi-Fi clients with Layer 2 connectivity (to each other) for supporting P2P communications, then such implementations can introduce security vulnerabilities in an IPv6 network such as clients sending out RA's or responding to DHCPv6 messages..

Solution*: Enable FHS solutions and force ALL P2P communications through the Subscriber Management Gateway*

FHS enablement would provide RA Guard - block all RA messages from all Wi-Fi clients, DHCPv6 Guard - recognize and block DHCPv6 responses from all clients, Source Guard - drop all traffic not sourced from client's IP address etc. Since the WLC responds to NS messages with entries in the local ND cache, it is possible to always force all Peer to Peer communication through the Gateway, which is the optimal place to enforce the policies and account for traffic usage.

Lawful / Legal Intercept (LI)

Cable MSOs are almost always required, as part of their regulatory compliance obligations, to keep a history of the IP addresses used by the subscribers, before placing any interception tap.

***Challenge**: Tracking and logging use of IPv6 addresses*

With SLAAC, IPv6 addresses are self assigned and each device can have (and use) multiple IPv6 addresses. Unfortunately, there is no centralized address management to log IPv6 addresses used by subscribers with SLAAC (and stateless DHCPv6).

***Solution**: Use out of band techniques to log and track IPv6 address usage*

One approach is to use the accounting start and stop records (RADIUS) that have client MAC addresses and IPv6 / IPv4 addresses and sometimes, even user names. These records will have to be generated as soon as a client gets assigned an IPv6 address and records will have to be stored in a database for historical reference / tracking for a period of time in accordance with regulatory compliance requirements.

There are other methods being discussed to address this issue. Please refer to this IETF draft - http://tools.ietf.org/html/draft-asati-dhc-ipv6-autoconfig-address-tracking-00

Prepaid Subscriptions

Prepaid subscriptions are offered based on time or volume or a combination of the two. All subscriber traffic and usage, regardless of IPv4 or IPv6 must be accounted for.

***Challenge**: Real time computation of usage and quota enforcement for dual-stack clients*

If there are two different sessions for IPv4 and IPv6 traffic, then accounting records will have to be reconciled (by the accounting/billing server) in real time to compute usage and enforce quotas. This can pose increasing difficulty to the accounting/billing server and sacrifice the accuracy of enforcement.

***Solution**: Enforce a single dual-stack session per device by having L2 connectivity between client and Subscriber Management Gateway*

With L2 connectivity between clients and Subscriber Management Gateway, session identification could be based on subscriber MAC address and all traffic to and from the subscriber (IPv4 and IPv6) will have the same identity making it easier to enforce volume or time based quotas real time. This has the added advantage of reducing the total number of sessions to be managed by the Subscriber Management Gateway.

Service Based Billing

MSOa may like the option of differentiated billing based on services consumed by the subscribers. Video, for example, could be billed at a different rate than voice. Similarly, access to MSO content could be provided at a reduced rate or tiered billing rate based on blocks of bandwidth consumption over a period of time.

***Challenge**: Consolidated accounting for services, regardless of whether they were consumed with IPv4 or IPv6*

Typically, services are independent of the type of stack (IPv4 or IPv6) used and in the case of dual-stack clients, the same services could be consumed using both IPv4 and IPv6. An example would be to identify a RTP stream for differentiated billing. RTP can run over both IPv4 and IPv6 because it contains no specific assumptions about the capabilities of the lower layers, except that they provide framing.

Solution: *Leverage QoS policies having both IPv4 and IPv6 traffic classes*

In Quality of Service design, Service flows (Services) are generally identified using characteristics from L4 through L7 and must include the option to match both IPv4 and IPv6 traffic classes. A service flow is set of traffic classes and is part of the services definition. When the services are installed for the subscriber on the Subscriber Management Gateway, an accounting start record, identifying the service name is sent. Interim accounting and accounting stop records are sent for all traffic that matches characteristics defined in all of the traffic classes associated with the service flow.

EVOLUTIONS

As we understand and assess the challenges, it becomes important to think outside the box and evolve IPv6 usage in SP Wi-Fi for further optimization. The below points discuss few such optimization opportunities:

1. If DAD could be disabled in case of a dedicated IPv6 prefix per subscriber device, then it would significantly reduce the control plane traffic on the Wi-Fi. Because the Gateway controls the prefix assignment, it can obviate the possibility of duplicate address assignment.

This would be helpful in the context of challenge described in section 4.6

2. If the device could acquire new IPv6 prefix/address after the move (from one AP (and WLC) to another AP (and WLC)) while keeping the old one(s), then tunneling of the control plane (i.e. original RA) and the data plane traffic could be limited to the existing IPv6 connections, while allowing the new IPv6 connections to use the new IPv6 prefix/address.

This would be helpful in the context of the challenge described in section 4.7, since the device would eventually stop using the old IPv6 prefix/addresses and could let go of them. This would cut down on having to tunnel the traffic from WLC to WLC and pave the way for optimal control plane and data plane traffic forwarding.

3. IP address logging is a critical prerequisite for Lawful Interception. When SLAAC is used for IPv6 address assignment, then IP address logging could fail. If the logging can still be done by the DHCP server (even if SLAAC is used), then it would be operationally beneficial to the MSOs. Thankfully, such a solution [4] is now discussed at the IETF DHC working group.

IPV6 STRATEGY / RECOMMENDATIONS

It is important to summarize the recommendations for IPv6 enabled SP Wi-Fi deployments:

1. Focus on IPv6 enablement for subscriber devices followed by infrastructure
2. Use Dual-stacking (i.e. IPv4 + IPv6) on subscriber devices
3. Leverage SLAAC for IPv6 address assignment (along with stateless DHCPv6) in short-term/medium-term and (stateful) DHCPv6 in long-term
4. Use a dedicated IPv6 prefix per subscriber device
5. Use one single IP session per subscriber device (whether dual-stacked or not)
6. Consider Unicast Router Advertisement
7. Preserve device's IP address during Mobility events, if possible
8. Use VRRP and VLANs on Subscriber Management Gateways for better high-availability and load-sharing
9. Prefer L2 connectivity between devices and Gateways (instead of multi-hop L3 connectivity)
10. Enforce First-Hop Security on WLC and Gateways, as well as subscribers' P2P communication via Gateways
11. Use HTTP redirection instead of L4 redirection
12. Leverage 'MAC address' as the IP session classifier
13. Log IPv6 addresses per subscriber.

SUMMARY

Is IPv6 over SP Wi-Fi a reality now? The answer is YES.

While there are obvious challenges to enabling IPv6 in SP Wi-Fi, there are solutions. There are certainly opportunities for optimization for further evolving IPv6 adoption in SP Wi-Fi networks. We urge the MSOs to start developing the strategy for enabling IPv6 in their SP Wi-Fi deployments (and pave the way for reducing the need for Carrier Grade NAT (CGN)44).

REFERENCES

[1] IEEE 802.11

[2] https://www.arin.net/resources/request/ipv4_countdown.html

[3] http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns673/data_sheet_c78-707266.html

[4] http://tools.ietf.org/html/draft-asati-dhc-ipv6-autoconfig-address-tracking-00

LEVERAGING OPENFLOW IN DOCSIS® NETWORKS

Chris Donley
CableLabs

Abstract

Networks are becoming virtualized. With the launch of new services and new demands on the network, operators are demanding greater flexibility, configuration consistency and control. One set of tools for meeting these demands is Software Defined Networking, specifically OpenFlow.

CableLabs, in partnership with MSOs and technology suppliers, has begun a technical exploration of how MSOs could leverage OpenFlow in a DOCSIS® environment. Our research is considering which subscriber services would see the greatest benefit from OpenFlow and how to architect OpenFlow into the DOCSIS network, specifically at the CMTS.

This paper will present findings from our research, specifically:

- *Targeted subscriber services enhanced through OpenFlow*
- *Key MSO benefits from the introduction of OpenFlow*
- *An architecture for hybrid Open Flow/DOCSIS networks*

INTRODUCTION

MSOs are expressing a growing interest in Software Defined Networking (SDN) and Network function Virtualization (NfV) technologies that have begun to emerge. These technologies promise a platform for rapid innovation and service deployment. They also promise a holistic view of the network - to be able to monitor and manage the network from a service perspective, rather than a device perspective. These trends, when fully realized, have the potential to improve operational efficiencies and accelerate the introduction of new services.

At CableLabs, we are working to enable the cable industry to capitalize on this paradigm shift towards software-controlled networks through:

- Knowledge sharing
- Architecture development & standards contributions
- Supplier readiness

This effort will contribute to a software-controlled network architecture allowing MSO resources to be configured, monitored, and optimized to improve subscriber experiences and business results.

Recently, we have analyzed MSO use cases to explore how OpenFlow, one SDN technology, could be applied in cable access networks, and the value to MSOs.

As we will discuss in this paper, our analysis found incremental enhancements to CMTS routing and forwarding for existing services such as L2/L3VPN configuration and deployment. OpenFlow also provides a step towards virtualization with traffic control features and management tools for managed firewalls and security, Carrier Grade NAT, and caching services.

This paper will also present an architecture for adding SDN/OpenFlow to cable access networks.

OPENFLOW

OpenFlow is one of the best-known SDN technologies. Developed by the Open Networking Foundation (ONF), OpenFlow specifies a two-way communication protocol by which a centralized OpenFlow controller (OFC) can add and remove "flow entries" (forwarding instructions) on network elements such as Ethernet switches and routers. This allows software programs interacting with the OpenFlow controller to programmatically control network forwarding.

OpenFlow ‘flows’ are logical constructs applied to traffic matching certain (and dynamically changing) classifiers. For instance, a flow could refer to a TCP connection to a particular website, all packets from a certain MAC or IP address, all packets tagged with a particular VLAN ID, or packets arriving on the same physical interface.

Each entry in the Flow Table contains three fields:

1. A classifier (packet header) that defines the flow;
2. An action, which directs the switch processing of the flow; and
3. Flow statistics, such as the number of packets or bytes for each flow.

OpenFlow defines three actions for each flow entry:

1. Forward the flow to a given port. This may also include instructions for manipulating the header or adding encapsulation headers;
2. Encapsulate the flow and send it to the OpenFlow controller. This is typically used for the first packet in a new flow; or
3. Drop this flow’s packets (e.g. to combat Denial of Service attacks).

While OpenFlow is an important technique for programmatic manipulation of traffic, it does not fully deliver the platform for rapid innovation and the holistic view of the network discussed above. In particular, it does not currently manage device provisioning; nor does it have a mechanism for directly controlling DOCSIS Quality of Service (QoS) techniques, bonding groups, or other DOCSIS-specific constructs.

OPENFLOW USE CASES

In order to identify the value provided by OpenFlow, it is helpful to analyze the protocol in the context of use cases based on services offered to cable subscribers. The use cases presented below assume all traffic forwarding is controlled by OpenFlow. As we will discuss in Section 0, it is likely that OpenFlow will first be added alongside traditional forwarding methods, rather than completely replacing traditional forwarding. However, analyzing the use cases below as pure OpenFlow use cases allows us to identify the value provided directly by OpenFlow.

In this paper, we will consider the following use cases:

- Basic Traffic Forwarding
- Traffic Optimization
- Security
- Virtual Private Networks
- Carrier Grade Network Address Translation
- Quality of Service (QoS)

Basic Traffic Forwarding

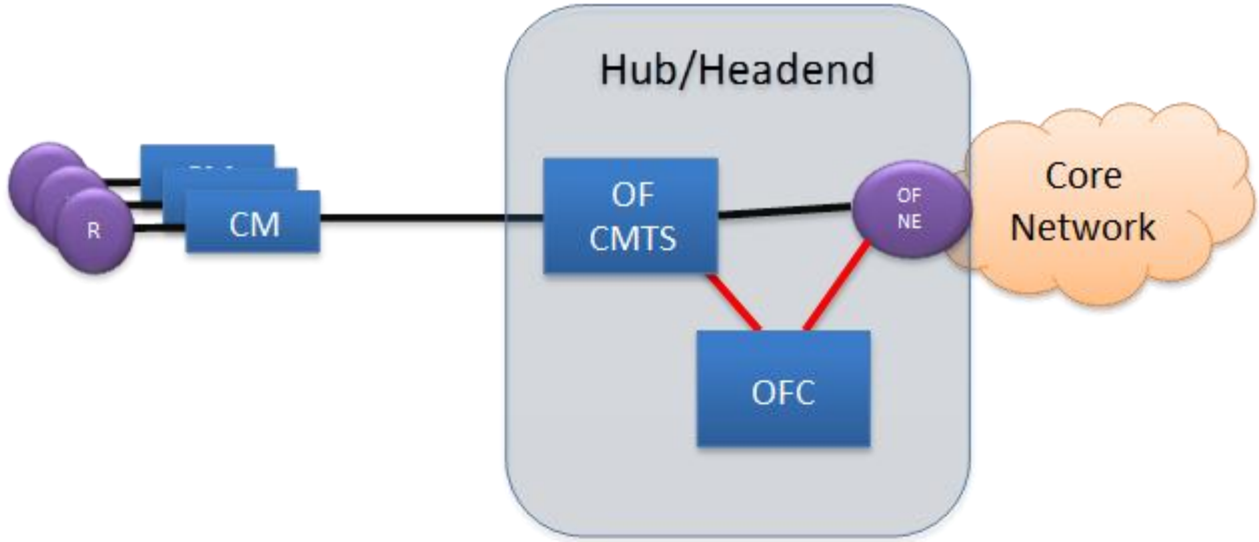

Figure 1: Basic Access Network with OpenFlow

In today's DOCSIS networks, the CMTS serves as a router. It uses routing protocols such as OSPF, ISIS, and BGP to learn the topology of the access network, serves as subscribers' default gateway, and routes traffic from subscriber devices to aggregation routers, which in turn route traffic across the MSO core network to the Internet or MSO servers.

In an OpenFlow environment, as shown in

As was described above, the initial http request for web content would be forwarded by the CMTS to the OpenFlow Controller. The OpenFlow Controller would invoke the content cache application to see if a cached copy is available. If the cache server has a copy of the requested content, it sets up a connection with the subscriber and serves the content locally. However, if the content cache does not have a copy available, or the content is too old, the OpenFlow Controller will direct the request to the content provider webserver. The response from the webserver can be sent directly to the client and also mirrored to the cache server for storage. Subsequent requests for this content would then be answered by the cache server. In the event of a cache server failure, the OpenFlow Controller would simply direct the request to the Internet.

Local caching offers several benefits to MSOs and their subscribers:

- Faster content retrieval for subscribers
- Reduced traffic on the MSO core network and transit links
- Increased content availability in the case of a webserver error or peak usage
- It facilitates the IPv4/IPv6 transition by serving as a proxy between the two protocols

OpenFlow facilitates the deployment of local cache servers by allowing MSOs to deploy the servers out-of-line, and direct specific flows to the caching servers. This would increase the overall scalability and reliability of the solution. It also allows more specific targeting of which content sites are cached, and which are not.

Security

Lawful Intercept

Lawful interception (LI) is a telecommunications function of collecting communications network data for a Law Enforcement Agency (LEA) for the purpose of analysis or evidence. Such data generally consists of signaling or network management information and in some instances, the content of the communications. These use cases explore how OpenFlow could facilitate lawful intercept functions.

Cable Broadband Intercept Specification (CBIS) Overview

The CableLabs CBIS specification identifies the specific interface points between the MSO and the LEA that has served the Broadband Intercept Order and enumerates the specific requirements for these interface points.

CBIS Outline

The following specific interfaces and logical functions have been identified and defined (as shown in Figure 4 and Figure 5) in order to meet the Law Enforcement's (LE) objectives and high-level requirements for Broadband Intercepts related to Transparency, Confidentiality, Authentication, Validation, Non-Repudiation, Correlation, Isolation, Completeness, Compression, and Encryption.

Access Function

The access function is a site-specific means of directing data to an out-of-band interface or to a packet stream interface. The access function may be implemented as an optical tap, a UDP data stream, a port mirror, or something else that is reliable and fast enough to manage multiple streams with no packet loss.

Mediation Function

The mediation function creates hashes and formats all events, headers and packet data depending on the type of intercept. The intercepts can be of two types: full packets or packet headers only. In either case, out of band data (e.g., DHCP) packets are captured. The mediation function has interfaces to collect raw data from the access function, and store formatted data at the broadband intercept function.

Broadband Intercept Function (BIF)

The broadband intercept function includes a buffer area that is used to store 24 hours of formatted data. The BIF is an optional function for MSOs; operators may choose to implement the BIF or request that the LEA provide the BIF. The operator needs to ensure that the buffer space is sufficiently sized on an LEA-by-LEA basis.

Collection Function

The collection function provides a secure means to deliver data to LEA. It is possible for more than one LEA to have simultaneous access to such data. Some LEAs will want to set up a VPN connection, while others will use SSH and/or portable storage. For each intercept, the operator and LEA must negotiate a single common solution for the topology and protocol (e.g., IPv4 or IPv6).

CBIS Operation

When a Lawful Intercept Order is received by the operator, the CBIS data collection begins. Access to CBIS equipment is strictly controlled by law and limited to prevent disclosure of the presence of an active intercept. CBIS data collection involves the following steps:

The operator identifies the cable modem associated with the subject facility.

CPE devices are identified via the cable modem MIB tables.

CBIS equipment is provisioned to capture the traffic, either by directly using CLI or using SNMP Tables such as extractions from the dot1dTpFdbTable.

The intercepted data is forwarded via the 'Access Function' to the 'Mediation Function'.

CBIS has enough information to create identification tags for expected data streams.

Data matching the IPv4 five-tuple or IPv6 six-tuple filters are formatted at the Mediation function and then passed to 'Broadband Intercept Function'. See Figure 5; the data includes both packet data and out of band messages.

At the end of the intercept, the data is forwarded to the 'Collection Function' for collection by the Law Enforcement Agency (LEA).

OpenFlow-Facilitated Lawful Intercept

As shown in Figure 5OpenFlow can facilitate Lawful Intercept, as all new traffic flows pass through the OpenFlow Controller. This would allow MSOs to initiate Lawful Intercept from the OpenFlow Controller, rather than the CMTS.

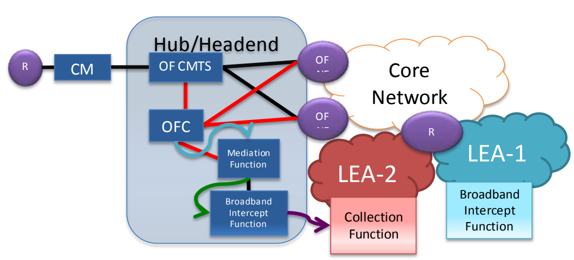

Figure 5: OpenFlow-facilitated Lawful Intercept

When a subscriber initiates a new flow, the CMTS forwards it to the OpenFlow Controller. The OpenFlow Controller will then invoke the Mediation Function to determine whether the flow is subject to an intercept order. If the Mediation Function finds that the flow subject to such an order, it directs the OpenFlow Controller to add a flow entry on the CMTS to mirrored the flow to one or more BIFs. The BIF then forwards data to the Collection Function, as today.

Using OpenFlow for Lawful Intercept provides several benefits to MSOs. First, it offers increased granularity in identifying flows, as it can look at additional fields beyond what is specified in CBIS. Second, it reduces the need to configure the CMTS for intercepts, preventing typographical errors from impacting CMTS operations. Third, it provides more control over the forwarding of “intercept flows” to one or more parties. Finally, it decouples CMTS code from Lawful Intercept updates. Should Lawful Intercept requirements change in the future, this approach allows them to be developed for a standalone application, and reduces the need for feature interaction testing on the CMTS.

Packet Inspection

OpenFlow can also be used to facilitate packet inspection. Since the OpenFlow controller only sees the first packet in every flow, it is still necessary to insert an inline Intrusion Detection System(IDS)/Intrusion Prevention System (IPS) to monitor traffic. When the IDS/IPS detects an intrusion, it notifies the OpenFlow Controller to help mitigate the attack. Depending on the nature of the attack, the OpenFlow Controller can instruct network elements to drop new and existing flows matching the attack signature or by setting per-flow bandwidth constraints. Similarly, connection-oriented attacks such as syn-floods can be mitigated by the OpenFlow Controller itself, without requiring an IPS. Using this approach, all security controls are added on an ad-hoc per-flow basis. This means that there is no configuration change on network elements to update after the attack abates.

This approach is more flexible than the systems available today, which directly update device configuration. Since the OpenFlow approach only temporarily updates flow tables, rather than issuing configuration commands, it is less likely to cause service problems on false-positives. Also, OpenFlow can augment IDS/IPS systems by directly mitigating connection-based Denial of Service attacks.

Managed firewall

With its ability to control per-flow forwarding behavior, OpenFlow brings MSOs the opportunity to offer a managed firewall service. The subscriber or MSO would first develop a security policy, defining which traffic should be forwarded and which should be dropped. Using OpenFlow, this policy is pushed out to both sides of the network via the OpenFlow Controller – customer edge (CMTS or CM) and aggregation/peering routers. When traffic from either direction enters the network, network elements forward the first packet to the OpenFlow Controller, which can check against the firewall policy and allow or deny traffic at the edge. The OpenFlow Controller can also maintain connection state, maintaining stateful firewall capabilities.

This approach provides a new managed service opportunity for MSOs without requiring dedicated equipment on the customer site. Also, services can be configured on the fly via a web portal, enabling self-provisioning, and reducing time and effort required for service changes. Also, the distributed nature of this firewall service allows for filtering close to both edges of the network, reducing transit bandwidth for malicious traffic within the MSO network.

Virtual Private Networks

Today, Layer 2 Virtual Private Networks (L2VPNs) are delivered over cable networks according to the CableLabs L2VPN specification. Provisioning is performed using a per-VPN CM config file that sets up L2VPN service flows and classifiers, and instructs the CMTS which encapsulation type to apply. The CM classifies upstream traffic flows onto L2VPN service flows, and the CMTS encapsulates the traffic using one of a number of encapsulation headers, including 802.1Q, 802.1ad, MPLS, L2TPv3, etc. While the specification allows for multipoint support, only point-to-point is implemented. L3VPN services are not specified for cable networks, although proprietary solutions are available.

As shown in **Error! Reference source not found.**, OpenFlow provides an alternative mechanism for delivering L2VPNs and L3VPNs, and can be implemented on either the CM or CMTS. If implemented on the CM, the CM receives upstream traffic and talks to the OpenFlow Controller about the new flow. The OpenFlow Controller checks with the VPN provisioning application and responds with encapsulation parameters. The CM then applies the encapsulation directly, in a manner similar to the model used by DPoE (DOCSIS provisioning of EPON). When the traffic reaches the CMTS, it checks with the OpenFlow Controller, and forwards the encapsulated traffic based on OpenFlow

Controller directions. When the CMTS receives encapsulated traffic to be sent downstream, it checks with the OpenFlow Controller and forwards the traffic to the CM. The CM then removes the encapsulation header based on OpenFlow flow entries and forwards the traffic to the destination. This approach does not require a per-CM config file, and requires minimal CMTS involvement with the VPN; however, as the encapsulation is applied at the CM, it could cause issues with large packets, as it would add headers that could cause the packet to exceed its MTU.

If OpenFlow is not enabled on the CM, a similar approach could be used for OpenFlow VPNs at the CMTS. In this case, the CMTS would talk to the OpenFlow Controller about a new upstream flow and receive encapsulation parameters, itself. It would then encapsulate and forward the upstream traffic as directed. Likewise, the CMTS would talk to the OpenFlow Controller about encapsulated downstream traffic, and then remove the encapsulation headers and forward the traffic to the RF interface, as directed.

This OpenFlow approach offers several advantages to MSOs. First, it allows dynamic provisioning with no per-CM config files. Second, it reduces technician touches of the CMTSs, as configuration is performed on an OpenFlow application. Third, it enables selective on-demand VPNs, e.g., for a home-office or road warrior. Finally, it enables multipoint support on the CMTS with no MAC address learning required at the CMTS, and no CMTS development to support the feature.

Performance Monitoring

Today's Metro Ethernet and L3VPN customers are demanding Service Level Agreements that define delay, loss, delay variation, and availability metrics. MSOs can monitor service performance against these metrics using ITU-T Y.1731 and MEF 35 performance monitoring standards. These standards define layer 2 messages to measure 1- and 2-way delay, loss, and delay variation. For true end-to-end measurements, they require support at each UNI (User-Network-Interface); however, current CMs do not implement this functionality. Using OpenFlow, it is possible to implement Y.1731 and MEF 35 without CM implementations.

When an MSO wants to measure service performance, a technician initiates the Performance Monitoring function via a network monitoring application. This application directs the OpenFlow Controller to send a special OpenFlow pkt_out message to the CM. This message tells the CM to generate a special message and forward it out its RF interface to the remote UNI. The remote UNI then responds; when the message reaches the CM, it sends it to the OpenFlow Controller, which sends data back to the Performance Monitoring application.

This approach offers a new feature that MSOs are requesting, but is not currently available. It also does not require Y.1731 support directly in the CM; only OpenFlow support, which also offers the benefits of the other use cases described here. This approach is also extensible, as new related features and new messages such as for Service Activation Testing would not require CM or CMTS development.

Carrier Grade Network Address Translation

Today, Carrier Grade NAT (CGN) is typically implemented in a dedicated hardware appliance or blade that serves 50,000 subscribers. OpenFlow can offer improvements to IPv4-IPv4 and IPv4-IPv6 Carrier Grade NAT by distributing the feature across multiple devices and bringing it closer to the subscriber.

As shown in Figure 8, CGN can be implemented as an OpenFlow application, rather than a standalone device. In this case, OpenFlow decouples upstream and downstream translations, and brings the feature closer to the subscriber.

In the upstream direction, the subscriber is provisioned with an IPv4 address in the 100.64.0.0/10 range and initiates a traffic flow. When the traffic reaches the CMTS, it forwards it to the OpenFlow Controller. The OpenFlow Controller communicates with the CGN application. For most flows, it then instructs the CMTS to rewrite the IPv4 source address to a predetermined CGN IPv4 address and the port to one selected using a deterministic algorithm. For flows requiring payload re-writes (e.g., VoIP), it instructs the CMTS to forward untranslated traffic to a dedicated network element with Application Layer Gateway (ALG) support to provide a better customer experience.

In the downstream direction, the process is reversed. An OpenFlow-enabled network element such as an aggregation router sends an incoming flow to the OpenFlow Controller, which then consults the CGN application, reverses the translation, and forwards the traffic to the CMTS.

This approach offers several benefits to MSOs. First, it allows deployment of the CGN feature closer to the subscriber. This lessens the latency and traffic engineering characteristic of a centralized approach. It also improves geolocation accuracy. Second, this approach does not require dedicated CGN hardware (except, potentially, for ALGs); however CGN software would be required. Third, this approach distributes translation duties across multiple devices, improving scalability. Finally, it provides an easier way than present methods for separating CGN subscribers from non-CGN subscribers.

In a similar manner, OpenFlow could also be used for IPv4-to-IPv6 NAT, a feature not available with current CGNs. This would allow IPv4-only clients (e.g. smart TVs) to talk to IPv6-only content servers without a proxy server, facilitating the transition to IPv6. This could also support inbound IPv4 services such as gaming and web hosting, if desired.

Quality of Service (QoS)

OpenFlow can be used to facilitate QoS, as shown in Figure 9. When the OpenFlow-enabled CMTS receives new subscriber-initiated traffic flows, it sends them to the OpenFlow Controller. For MSO services such as VoIP, the OpenFlow Controller directs the CMTS to set the Differentiated Services Code Point (DSCP) and 802.1p bits of the traffic flow. As traffic progresses through the network, network elements apply QoS policies as determined by the DSCP and 802.1p bits.

OpenFlow does not have a way to directly influence DOCSIS QoS on the RF network. However, when DOCSIS QoS is needed, the OpenFlow Controller can send the traffic flow to the PacketCable MultiMedia (PCMM) Policy Server (PS), which then communicates with the CMTS and initiates DOCSIS QoS. This helps with PCMM deployment by only exposing the PS to traffic flows for which DOCSIS QoS is required.

For third-party services, OpenFlow can also help direct traffic to the nearest server using Global Server Load Balancing (GSLB). Usually, content providers use GSLB to monitor the path from their servers to subscribers and direct subscribers to the closest server; in this scenario, the MSO uses GSLB to direct subscriber traffic according to its criteria, such as the lowest latency path, lowest cost path, etc. As before, the CMTS sends new traffic flows to the OpenFlow Controller. In this case, the OpenFlow Controller invokes the GSLB application, which measures connectivity to the content provider's data centers and identifies the best path. The OpenFlow Controller then directs the CMTS to forward the subscriber's new flow to the appropriate data center according to its criteria. If necessary, the OpenFlow Controller could instruct the CMTS to rewrite the packet headers to reach the appropriate server.

This approach offers MSOs increased flexibility in the deployment of QoS. For MSO-provided services, OpenFlow sets QoS bits on traffic in the access and core network. While it cannot directly enable DOCSIS QoS, OpenFlow can interface with PCMM. This allows for more dynamic control of QOS, compared to today's deployments. OpenFlow also offers the ability to optimize the path between the subscriber and content server(s).

HYBRID OPENFLOW CMTS ARCHITECTURE

At CableLabs, we expect that MSOs will not move directly from traditional networking approaches to SDN, but rather will phase it in over time. As such, the architecture presented herein allows MSOs to leverage SDN in the access network for service agility, while preserving existing operational models for some services.

It is important to note that the architecture described below is the view of the author, and is not incorporated into any CableLabs specification, requirements document, or technical report.

The Big Picture

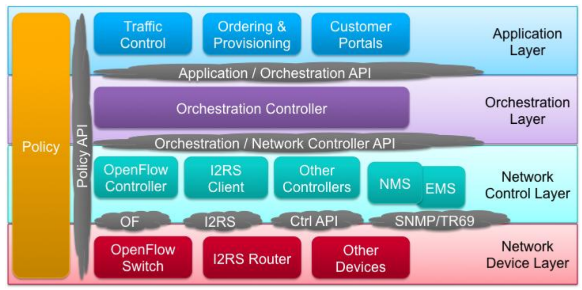

Figure 10: Emerging Network Architecture, the CMTS serves as a Layer 2 device (although, as we will discuss below, it can appear to exhibit Layer 3 behavior). When a new upstream traffic flow is initiated, the CMTS sends the first one or two packets to the OpenFlow controller, which determines where and how to forward the packet. The OpenFlow controller directs the CMTS to add a flow entry into its Flow Table, and the CMTS makes future forwarding decisions based on the flow entry. The OpenFlow Controller could also install a wildcard flow entry, matching multiple traffic types and/or destinations, to speed up the processing of aggregated flows.

This change to the CMTS' forwarding behavior would be transparent to subscribers. Subscriber devices would still use DHCP for provisioning. Devices supporting IPv4 would use DHCPv4 to acquire their default gateway router address, while devices supporting IPv6 would receive Router Advertisement (RA) messages informing them of their default router. Such devices would send traffic to the CMTS as they do today, with no knowledge that the CMTS uses OpenFlow. The only difference is that instead of using the CMTS' IP address as the default gateway, subscriber devices would use the IP address of the aggregation router, instead.

Traffic Optimization

DNS Caching

The Domain Name System (DNS) allows for caching as a means to attain higher scalability by storing DNS information locally, instead of sending all queries to a central server. Indeed, MSOs have deployed caching name servers for years. Adding OpenFlow in the access network helps to push the DNS infrastructure closer to the subscribers, reducing demands on the core network while providing more responsive service to subscribers.

would receive the queries and direct them to the OpenFlow Controller. The OpenFlow Controller would then invoke a DNS caching application to see if it can respond to the query. If the caching application can respond, it will send a DNS response back to the subscriber. If the DNS caching application cannot respond (e.g., because it is unavailable) or if it does not have current information, the OpenFlow Controller forwards the query to the MSO's centralized DNS server. When the centralized DNS server responds, the OpenFlow Controller sends one copy to the subscriber and a second copy to the caching application.

Local DNS caching provides several benefits:

- It reduces DNS traffic on the MSO core network
- It offers faster response times for subscribers

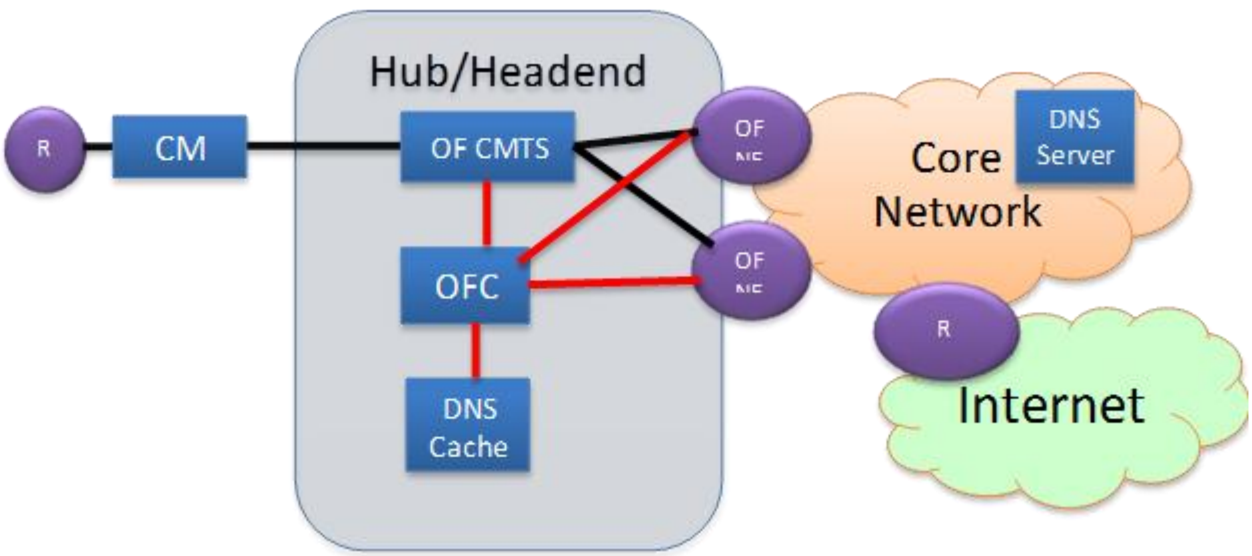

Figure 2: OpenFlow-facilitated DNS caching

OpenFlow provides additional control for serving DNS requests, as shown in Figure 2. Subscribers would generate DNS queries for services they use (e.g., www.google.com) and send them to the MSO's anycast DNS server address, as they do today. The CMTS

- It opens the possibility of tailoring the DNS response for local needs. For instance, MSOs could implement Global Server Load Balancing as part of the caching application, where the DNS server measures availability and latency of Internet content servers and directs local subscribers to the server with the least latency for them.

OpenFlow facilitates local DNS caching by allowing MSOs to direct traffic to local servers. Without OpenFlow, MSOs use anycast addressing to direct traffic to the nearest server. This requires more complicated routing and can cause localized outages during server failures before the anycast route is withdrawn. Using OpenFlow simplifies the deployment of DNS caching servers in local networks, and can reduce the impact of local failures by forwarding a copy of the request to core servers.

Content Caching

As with DNS, content caching has been available to MSOs for some time. However, content cache servers are generally required to be deployed in-line to be able to capture and respond to http requests. They typically have to examine all traffic, whether it can be served by the cache servers or not. Also, MSOs deploying content caches need to plan for high-availability in the event of a server failure. OpenFlow facilitates the deployment of caching servers, as shown in Figure 3.

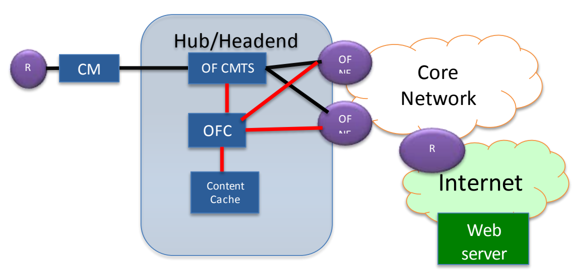

Figure 3: OpenFlow-facilitated Content Caching

As was described above, the initial http request for web content would be forwarded by the CMTS to the OpenFlow Controller. The OpenFlow Controller would invoke the content cache application to see if a cached copy is available. If the cache server has a copy of the requested content, it sets up a connection with the subscriber and serves the content locally. However, if the content cache does not have a copy available, or the content is too old, the OpenFlow Controller will direct the request to the content provider webserver. The response from the webserver can be sent directly to the client and also mirrored to the cache server for storage. Subsequent requests for this content would then be answered by the cache server. In the event of a cache server failure, the OpenFlow Controller would simply direct the request to the Internet.

Local caching offers several benefits to MSOs and their subscribers:

- Faster content retrieval for subscribers
- Reduced traffic on the MSO core network and transit links
- Increased content availability in the case of a webserver error or peak usage
- It facilitates the IPv4/IPv6 transition by serving as a proxy between the two protocols

OpenFlow facilitates the deployment of local cache servers by allowing MSOs to deploy the servers out-of-line, and direct specific flows to the caching servers. This would increase the overall scalability and reliability of the solution. It also allows more specific targeting of which content sites are cached, and which are not.

Security

Lawful Intercept

Lawful interception (LI) is a telecommunications function of collecting communications network data for a Law Enforcement Agency (LEA) for the purpose of analysis or evidence. Such data generally consists of signaling or network management information and in some instances, the content of the communications. These use cases explore how OpenFlow could facilitate lawful intercept functions.

Cable Broadband Intercept Specification (CBIS) Overview

The CableLabs CBIS specification identifies the specific interface points between the MSO and the LEA that has served the Broadband Intercept Order and enumerates the specific requirements for these interface points.

CBIS Outline

The following specific interfaces and logical functions have been identified and defined (as shown in Figure 4 and Figure 5) in order to meet the Law Enforcement's (LE) objectives and high-level requirements for Broadband Intercepts related to Transparency, Confidentiality, Authentication, Validation, Non-Repudiation, Correlation, Isolation, Completeness, Compression, and Encryption.

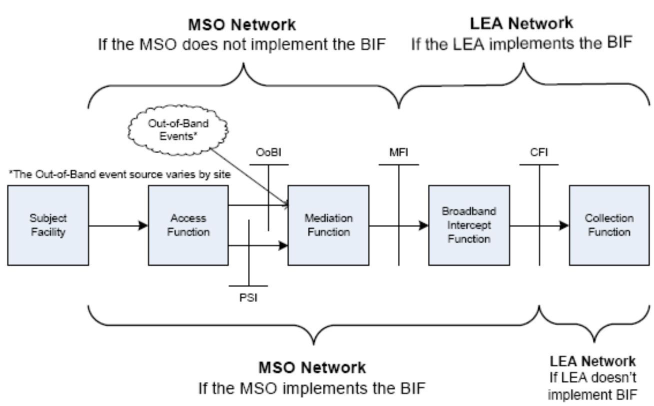

Figure 4: CBIS Broadband Intercept Interfaces

Access Function

The access function is a site-specific means of directing data to an out-of-band interface or to a packet stream interface. The access function may be implemented as an optical tap, a UDP data stream, a port mirror, or something else that is reliable and fast enough to manage multiple streams with no packet loss.

Mediation Function

The mediation function creates hashes and formats all events, headers and packet data depending on the type of intercept. The intercepts can be of two types: full packets or packet headers only. In either case, out of band data (e.g., DHCP) packets are captured. The mediation function has interfaces to collect raw data from the access function, and store formatted data at the broadband intercept function.

Broadband Intercept Function (BIF)

The broadband intercept function includes a buffer area that is used to store 24 hours of formatted data. The BIF is an optional function for MSOs; operators may choose to implement the BIF or request that the LEA provide the BIF. The operator needs to ensure that the buffer space is sufficiently sized on an LEA-by-LEA basis.

Collection Function

The collection function provides a secure means to deliver data to LEA. It is possible for more than one LEA to have simultaneous access to such data. Some LEAs will want to set up a VPN connection, while others will use SSH and/or portable storage. For each intercept, the operator and LEA must negotiate a single common solution for the topology and protocol (e.g., IPv4 or IPv6).

CBIS Operation

When a Lawful Intercept Order is received by the operator, the CBIS data collection begins. Access to CBIS equipment is strictly controlled by law and limited to prevent disclosure of the presence of an active intercept. CBIS data collection involves the following steps:

The operator identifies the cable modem associated with the subject facility.

CPE devices are identified via the cable modem MIB tables.

CBIS equipment is provisioned to capture the traffic, either by directly using CLI or using SNMP Tables such as extractions from the dot1dTpFdbTable.

The intercepted data is forwarded via the 'Access Function' to the 'Mediation Function'.

CBIS has enough information to create identification tags for expected data streams.

Data matching the IPv4 five-tuple or IPv6 six-tuple filters are formatted at the Mediation function and then passed to 'Broadband Intercept Function'. See Figure 5; the data includes both packet data and out of band messages.

At the end of the intercept, the data is forwarded to the 'Collection Function' for collection by the Law Enforcement Agency (LEA).

OpenFlow-Facilitated Lawful Intercept

As shown in Figure 5OpenFlow can facilitate Lawful Intercept, as all new traffic flows pass through the OpenFlow Controller. This would allow MSOs to initiate Lawful Intercept from the OpenFlow Controller, rather than the CMTS.

Figure 6: CBIS Logical Network

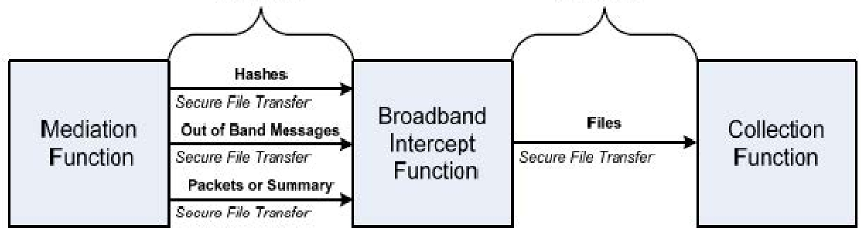

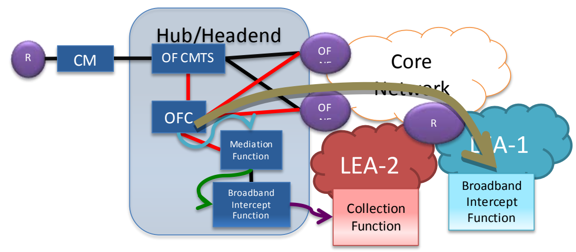

Figure 5: OpenFlow-facilitated Lawful Intercept

When a subscriber initiates a new flow, the CMTS forwards it to the OpenFlow Controller. The OpenFlow Controller will then invoke the Mediation Function to determine whether the flow is subject to an intercept order. If the Mediation Function finds that the flow subject to such an order, it directs the OpenFlow Controller to add a flow entry on the CMTS to mirrored the flow to one or more BIFs. The BIF then forwards data to the Collection Function, as today.

Using OpenFlow for Lawful Intercept provides several benefits to MSOs. First, it offers increased granularity in identifying flows, as it can look at additional fields beyond what is specified in CBIS. Second, it reduces the need to configure the CMTS for intercepts, preventing typographical errors from impacting CMTS operations. Third, it provides more control over the forwarding of “intercept flows” to one or more parties. Finally, it decouples CMTS code from Lawful Intercept updates. Should Lawful Intercept requirements change in the future, this approach allows them to be developed for a standalone application, and reduces the need for feature interaction testing on the CMTS.

Packet Inspection

OpenFlow can also be used to facilitate packet inspection. Since the OpenFlow controller only sees the first packet in every flow, it is still necessary to insert an inline Intrusion Detection System(IDS)/Intrusion Prevention System (IPS) to monitor traffic. When the IDS/IPS detects an intrusion, it notifies the OpenFlow Controller to help mitigate the attack. Depending on the nature of the attack, the OpenFlow Controller can instruct network elements to drop new and existing flows matching the attack signature or by setting per-flow bandwidth constraints. Similarly, connection-oriented attacks such as syn-floods can be mitigated by the OpenFlow Controller itself, without requiring an IPS. Using this approach, all security controls are added on an ad-hoc per-flow basis. This means that there is no configuration change on network elements to update after the attack abates.

This approach is more flexible than the systems available today, which directly update device configuration. Since the OpenFlow approach only temporarily updates flow tables, rather than issuing configuration commands, it is less likely to cause service problems on false-positives. Also, OpenFlow can augment IDS/IPS systems by directly mitigating connection-based Denial of Service attacks.

Managed firewall

With its ability to control per-flow forwarding behavior, OpenFlow brings MSOs the opportunity to offer a managed firewall service. The subscriber or MSO would first develop a security policy, defining which traffic should be forwarded and which should be dropped. Using OpenFlow, this policy is pushed out to both sides of the network via the OpenFlow Controller – customer edge (CMTS or CM) and aggregation/peering routers. When traffic from either direction enters the network, network elements forward the first packet to the OpenFlow Controller, which can check against the firewall policy and allow or deny traffic at the edge. The OpenFlow Controller can also maintain connection state, maintaining stateful firewall capabilities.

This approach provides a new managed service opportunity for MSOs without requiring dedicated equipment on the customer site. Also, services can be configured on the fly via a web portal, enabling self-provisioning, and reducing time and effort required for service changes. Also, the distributed nature of this firewall service allows for filtering close to both edges of the network, reducing transit bandwidth for malicious traffic within the MSO network.

Virtual Private Networks

Today, Layer 2 Virtual Private Networks (L2VPNs) are delivered over cable networks according to the CableLabs L2VPN specification. Provisioning is performed using a per-VPN CM config file that sets up L2VPN service flows and classifiers, and instructs the CMTS which encapsulation type to apply. The CM classifies upstream traffic flows onto L2VPN service flows, and the CMTS encapsulates the traffic using one of a number of encapsulation headers, including 802.1Q, 802.1ad, MPLS, L2TPv3, etc. While the specification allows for multipoint support, only point-to-point is implemented. L3VPN services are not specified for cable networks, although proprietary solutions are available.

As shown in **Error! Reference source not found.**, OpenFlow provides an alternative mechanism for delivering L2VPNs and L3VPNs, and can be implemented on either the CM or CMTS. If implemented on the CM, the CM receives upstream traffic and talks to the OpenFlow Controller about the new flow. The OpenFlow Controller checks with the VPN provisioning application and responds with encapsulation parameters. The CM then applies the encapsulation directly, in a manner similar to the model used by DPoE (DOCSIS provisioning of EPON). When the traffic reaches the CMTS, it checks with the OpenFlow Controller, and forwards the encapsulated traffic based on OpenFlow Controller directions. When the CMTS receives encapsulated traffic to be sent downstream, it checks with the OpenFlow Controller and forwards the traffic to the CM. The CM then removes the encapsulation header based on OpenFlow flow entries and forwards the traffic to the destination. This approach does not require a per-CM config file, and requires minimal CMTS involvement with the VPN; however, as the encapsulation is applied at the CM, it could cause issues with large packets, as it would add headers that could cause the packet to exceed its MTU.

If OpenFlow is not enabled on the CM, a similar approach could be used for OpenFlow VPNs at the CMTS. In this case, the CMTS would talk to the OpenFlow Controller about a new upstream flow and receive encapsulation parameters, itself. It would then encapsulate and forward the upstream traffic as directed. Likewise, the CMTS would talk to the OpenFlow Controller about encapsulated downstream traffic, and then remove the encapsulation headers and forward the traffic to the RF interface, as directed.

This OpenFlow approach offers several advantages to MSOs. First, it allows dynamic provisioning with no per-CM config files. Second, it reduces technician touches of the CMTSs, as configuration is performed on an OpenFlow application. Third, it enables selective on-demand VPNs, e.g., for a home-office or road warrior. Finally, it enables multipoint support on the CMTS with no MAC address learning required at the CMTS, and no CMTS development to support the feature.

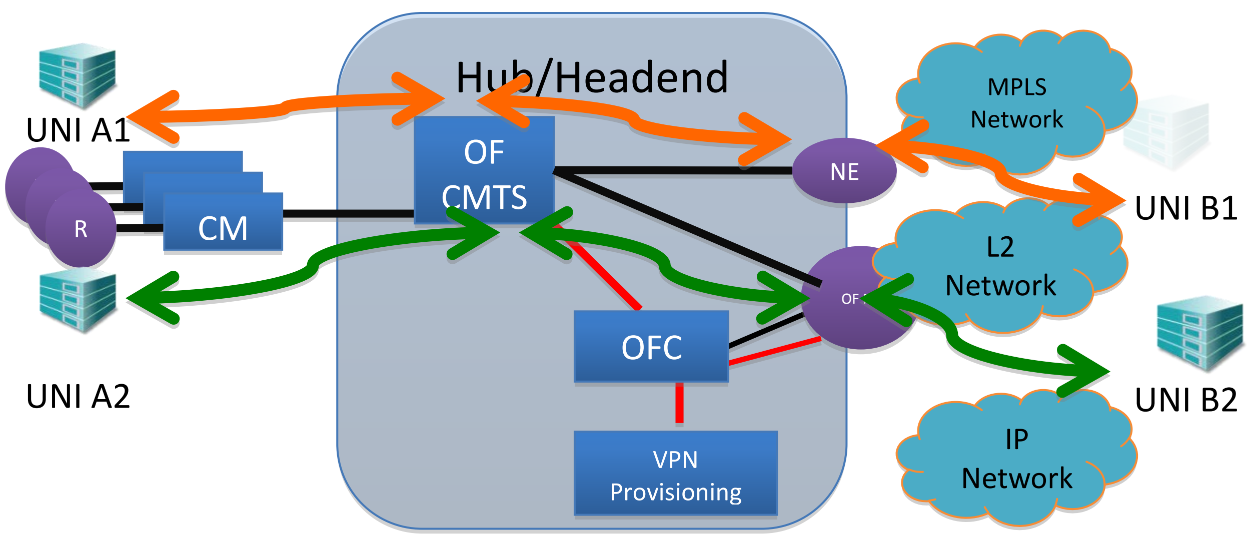

Figure 7: OpenFlow-facilitated VPNs

Performance Monitoring

Today's Metro Ethernet and L3VPN customers are demanding Service Level Agreements that define delay, loss, delay variation, and availability metrics. MSOs can monitor service performance against these metrics using ITU-T Y.1731 and MEF 35 performance monitoring standards. These standards define layer 2 messages to measure 1- and 2-way delay, loss, and delay variation. For true end-to-end measurements, they require support at each UNI (User-Network-Interface); however, current CMs do not implement this functionality. Using OpenFlow, it is possible to implement Y.1731 and MEF 35 without CM implementations.

When an MSO wants to measure service performance, a technician initiates the Performance Monitoring function via a network monitoring application. This application directs the OpenFlow Controller to send a special OpenFlow pkt_out message to the CM. This message tells the CM to generate a special message and forward it out its RF interface to the remote UNI. The remote UNI then responds; when the message reaches the CM, it sends it to the OpenFlow Controller, which sends data back to the Performance Monitoring application.

This approach offers a new feature that MSOs are requesting, but is not currently available. It also does not require Y.1731 support directly in the CM; only OpenFlow support, which also offers the benefits of the other use cases described here. This approach is also extensible, as new related features and new messages such as for Service Activation Testing would not require CM or CMTS development.

<u>Carrier Grade Network Address Translation</u>

Today, Carrier Grade NAT (CGN) is typically implemented in a dedicated hardware appliance or blade that serves 50,000 subscribers. OpenFlow can offer improvements to IPv4-IPv4 and IPv4-IPv6 Carrier Grade NAT by distributing the feature across multiple devices and bringing it closer to the subscriber.

As shown in Figure 8, CGN can be implemented as an OpenFlow application, rather than a standalone device. In this case,

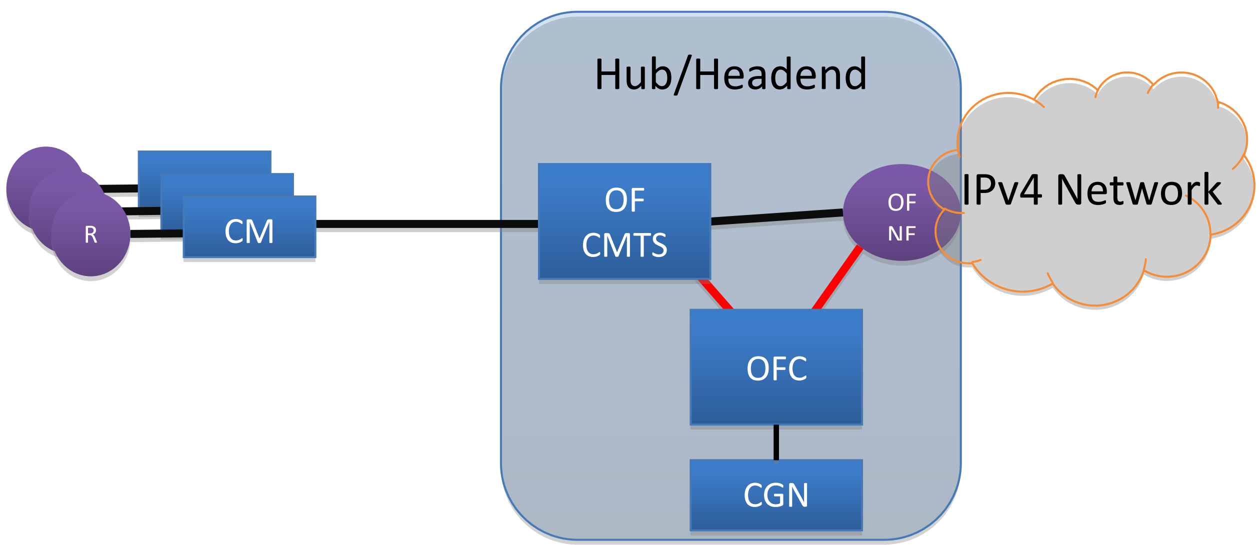

Figure 8: Carrier Grade NAT using OpenFlow

OpenFlow decouples upstream and downstream translations, and brings the feature closer to the subscriber.

In the upstream direction, the subscriber is provisioned with an IPv4 address in the 100.64.0.0/10 range and initiates a traffic flow. When the traffic reaches the CMTS, it forwards it to the OpenFlow Controller. The OpenFlow Controller communicates with the CGN application. For most flows, it then instructs the CMTS to rewrite the IPv4 source address to a predetermined CGN IPv4 address and the port to one selected using a deterministic algorithm. For flows requiring payload re-writes (e.g., VoIP), it instructs the CMTS to forward untranslated traffic to a dedicated network element with Application Layer Gateway (ALG) support to provide a better customer experience.

In the downstream direction, the process is reversed. An OpenFlow-enabled network element such as an aggregation router sends an incoming flow to the OpenFlow Controller, which then consults the CGN application, reverses the translation, and forwards the traffic to the CMTS.

This approach offers several benefits to MSOs. First, it allows deployment of the CGN feature closer to the subscriber. This lessens the latency and traffic engineering characteristic of a centralized approach. It also improves geolocation accuracy. Second, this approach does not require dedicated CGN hardware (except, potentially, for ALGs); however CGN software would be required. Third, this approach distributes translation duties across multiple devices, improving scalability. Finally, it provides an easier way than present methods for separating CGN subscribers from non-CGN subscribers.

In a similar manner, OpenFlow could also be used for IPv4-to-IPv6 NAT, a feature not available with current CGNs. This would allow IPv4-only clients (e.g. smart TVs) to talk to IPv6-only content servers without a proxy server, facilitating the transition to IPv6. This could also support inbound IPv4 services such as gaming and web hosting, if desired.

Quality of Service (QoS)

OpenFlow can be used to facilitate QoS, as shown in Figure 9. When the OpenFlow-enabled CMTS receives new subscriber-initiated traffic flows, it sends them to the OpenFlow Controller. For MSO services such as VoIP, the OpenFlow Controller directs the CMTS to set the Differentiated Services Code Point (DSCP) and 802.1p bits of the traffic flow. As traffic progresses through the network, network elements apply QoS policies as determined by the DSCP and 802.1p bits.

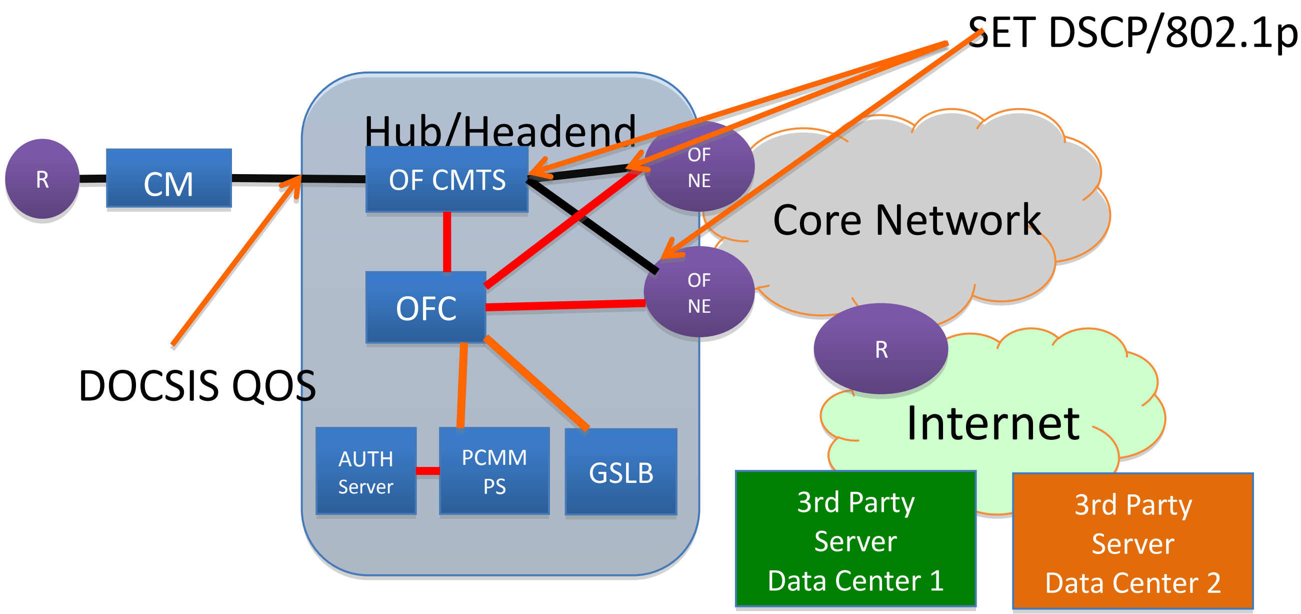

Figure 9: OpenFlow-facilitated QoS

OpenFlow does not have a way to directly influence DOCSIS QoS on the RF network. However, when DOCSIS QoS is needed, the OpenFlow Controller can send the traffic flow to the PacketCable MultiMedia (PCMM) Policy Server (PS), which then communicates with the CMTS and initiates DOCSIS QoS. This helps with PCMM deployment by only exposing the PS to traffic flows for which DOCSIS QoS is required.

For third-party services, OpenFlow can also help direct traffic to the nearest server using Global Server Load Balancing (GSLB). Usually, content providers use GSLB to monitor the path from their servers to subscribers and direct subscribers to the closest server; in this scenario, the MSO uses GSLB to direct subscriber traffic according to its criteria, such as the lowest latency path, lowest cost path, etc. As before, the CMTS sends new traffic flows to the OpenFlow Controller. In this case, the OpenFlow Controller invokes the GSLB application, which measures connectivity to the content provider's data centers and identifies the best path. The OpenFlow Controller then directs the CMTS to forward the subscriber's new flow to the appropriate data center according to its criteria. If necessary, the OpenFlow Controller could instruct the CMTS to rewrite the packet headers to reach the appropriate server.

This approach offers MSOs increased flexibility in the deployment of QoS. For MSO-provided services, OpenFlow sets QoS bits on traffic in the access and core network. While it cannot directly enable DOCSIS QoS, OpenFlow can interface with

PCMM. This allows for more dynamic control of QOS, compared to today's deployments. OpenFlow also offers the ability to optimize the path between the subscriber and content server(s).

HYBRID OPENFLOW CMTS ARCHITECTURE

At CableLabs, we expect that MSOs will not move directly from traditional networking approaches to SDN, but rather will phase it in over time. As such, the architecture presented herein allows MSOs to leverage SDN in the access network for service agility, while preserving existing operational models for some services.

It is important to note that the architecture described below is the view of the author, and is not incorporated into any CableLabs specification, requirements document, or technical report.

The Big Picture

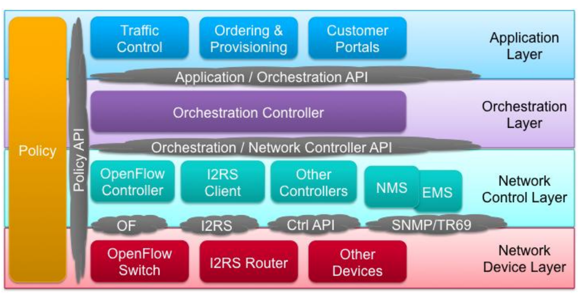

Figure 10: Emerging Network Architecture

While this paper has discussed the applicability of OpenFlow up to this point, it is important to note that OpenFlow by itself does not fulfill the promise of SDN. As shown in Figure 10, OpenFlow is one component of a larger architecture that will provide MSOs with the service agility and holistic control they will need in coming years. Other technologies are being developed to fill out the capabilities of this new architecture and to interface between new applications and different legacy and emerging network technologies.

OpenFlow is perhaps the most developed and widely researched SDN technology, however. Therefore, the next few sections describe an architecture for adding OpenFlow to the CMTS in a manner that also allows for addition of additional SDN technologies as they become available and relevant to MSOs. It deals specifically with traffic differentiation between OpenFlow and non-OpenFlow traffic and the OpenFlow forwarding model to use for the CMTS. This paper does not address topics such as redundancy or feature migration.

Traffic Differentiation

The first questions to answer is how to differentiate upstream traffic forwarding to be directed by OpenFlow from traffic to be forwarded using conventional means. There are five possible approaches to traffic segmentation at the CMTS:

1. Separate DOCSIS channels – establish separate pools of bonded RF channels for OF and non-OF traffic. Use the CM config file to direct traffic to a particular channel.
2. Separate DOCSIS Service Flows – establish a separate Service Flow for OF traffic, and use DOCSIS classifiers to direct traffic into the OF Service Flow. The CMTS processes any traffic received on an OF Service Flow using OpenFlow, and all other traffic using traditional forwarding methods.
3. Per traffic type – configure an access list in the CMTS that segments traffic by destination port, with some well-known ports processed using OpenFlow and others using traditional methods.
4. Per Source IP or MAC address – configure an access list in the CMTS

that segments traffic by source IP or MAC address, with traffic from predetermined source addresses processed using OpenFlow, and all other traffic processed using traditional methods.

5. Sequentially – the CMTS checks the OpenFlow Flow Table first, then the CMTS FIB second. This means that all traffic flows are first sent to the OpenFlow controller; if the OpenFlow Controller sets up a Flow Table entry, additional traffic from that flow is processed

We believe that reserving separate DOCSIS RF channels for OpenFlow traffic is economically infeasible. Also, sequential processing increases latency, as every traditionally-managed flow would need to be processed by OpenFlow first. Therefore, we concentrate our analysis on applying the use cases described above to the remaining three approaches, as shown in

. Approaches well-matched to the use cases receive a (✓); approaches partially matched to the use case, or that can only work in limited circumstances receive a (~); and approaches unsuitable for the use case receive a (✗).

Table 1: Viability of traffic classification methods

Use Case	*Traffic Type*	*Per-SF*	*Per-Port*	*Per-IP/MAC*
DNS Caching	DNS	✓	✓	✗
Content Caching	Varies; primarily http, https	✓	✓	✗
Lawful Intercept	All	✗	✗	✓
IDS/IPS	All	✓	~	✓
Managed Firewall	All	✓	~	~
L2VPN/L3VPN	All	✓	✗	~
Carrier Grade NAT	All Primarily http/https and DNS	✓	~	✓
QoS	VoIP: SIP and RTP Video: http Gaming and other OTT services: varies	✓	~	~

As shown above, per-service flow separation of OpenFlow and traditionally managed traffic appears to fit most use cases, except Lawful Intercept. As service flow establishment would send a message to the CM that could be observed by a subject under an intercept order, alternative approaches (particularly per-source MAC or IP address) would be required to initiate Lawful Intercept without notifying the subject.

One way to take advantage of service flow-based classification, while preserving flexibility for additional use cases such as Lawful Intercept, is a hybrid approach shown in **Error! Reference source not found.**. Traffic arriving on service flows identified as OpenFlow service flows are processed directly by the OpenFlow flow table. Traffic arriving on other service flows can be sent to an Access Control List (ACL), which directs specific flows to the OpenFlow flow table and all remaining traffic to the CMTS Forwarding Information Base (FIB). This approach satisfies all of the use cases described above, and offers MSOs significant flexibility to experiment with OpenFlow in the lab and field trials.

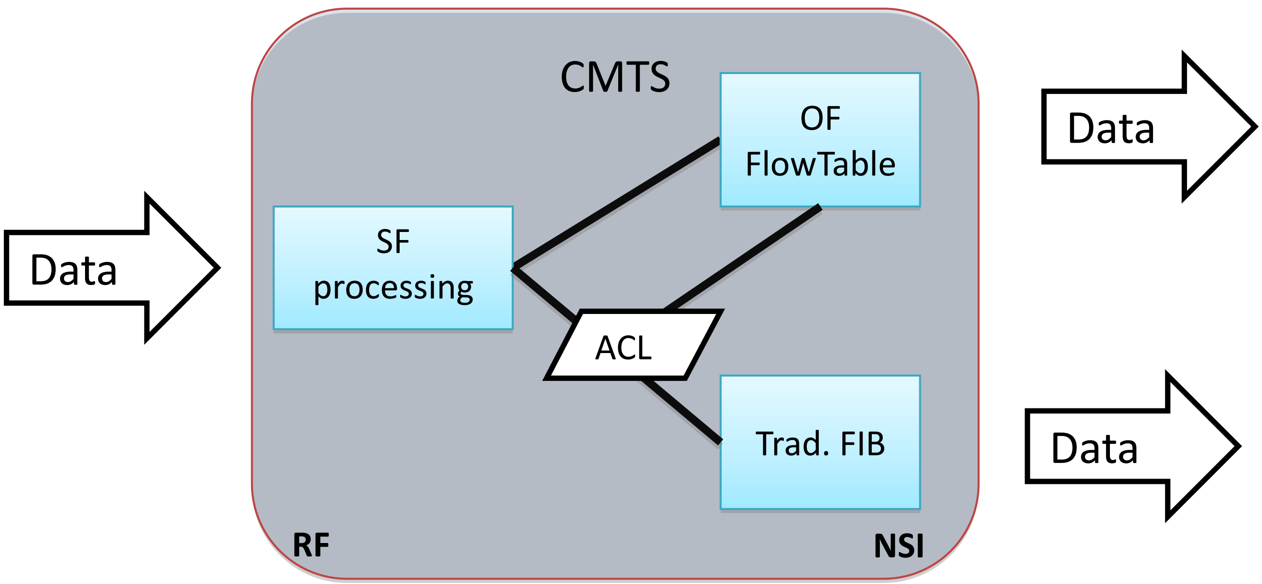

Figure 11: Hybrid SF-based Traffic Separation

<u>Traffic Forwarding</u>

CableLabs identified two hybrid OpenFlow CMTS models, one referred to as the “L2 Model”, and the second as the “L2/L3 Model”. Both are described below.

L2 Model

In the L2 model, the CMTS behaves solely as a Layer 2 device. Within a headend or hub site, there is a common Layer 2 domain for all DOCSIS interfaces on all CMTSs. Subscriber devices are provisioned from the same IPv4 subnet/IPv6 prefix, and they receive the address of the aggregation router (not the CMTS) as their default gateway.

When subscriber traffic arrives at the CMTS, the CMTS talks to the OpenFlow Controller and installs flow entries based on load balancing, failover, traffic control, premium services, and other factors that tell it where to forward traffic. In order to reduce MAC learning on the routers, the CMTS transforms the Ethernet header on upstream flows to use its source MAC address, rather than the subscriber device. In order to keep broadcast and multicast traffic rates low in the access network, the OpenFlow Controller can either direct IPv6 Neighbor Discovery (ND) and IPv4 ARP messages directly to their target nodes, without flooding the network, or suppress them on the access network and respond with pkt_out messages directing the CMTS to generate the messages locally.

In the downstream direction, routers add ARP and ND entries mapping the subscriber IP address to the CMTS MAC address. Routes to subscriber IPv6 prefixes would be mapped to the appropriate customer router, and traffic directed to the respective CMTS. As in the upstream direction, the CMTS would remap downstream flows to point to the subscriber MAC address.

This approach offers MSOs several benefits. First, it allows MSOs more granular control over traffic forwarding in the access network for load balancing, failover, and traffic control. It also offers the possibility of separate paths for premium services. Second, OpenFlow provides operational benefits such as eliminating the need for IP address renumbering during node splits and reducing the need for routing protocols in the CMTS.

L2/L3 Model

In the L2/L3 model, the CMTS behaves like a Layer 3 device, as it does today. During provisioning, each CMTS is assigned a different subnet for each RF interface. Subscriber devices are then provisioned to use the CMTS as the default gateway router. As it does today, the CMTS would be responsible for sending IPv6 Router Advertisement and IPv4 ARP messages. However, the CMTS makes traffic forwarding decisions based on OpenFlow, rather than traditional methods.

When the CMTS receives a new subscriber flow, it talks to the OpenFlow Controller to learn where to direct the flow and how to transform Ethernet headers. In this case, the transformation looks like the routing process. The CMTS changes the Ethernet source and destination addresses from Subscriber MAC:CMTS MAC to CMTS MAC:Router MAC, where the router MAC is the MAC address of the router selected by the OpenFlow Controller as the next hop. As discussed above, router selection could be based on load balancing, failover, traffic control, and premium services.

In the downstream direction, routers can either use traditional routing or OpenFlow to direct traffic to the correct CMTS for forwarding to subscribers. When the flow reaches the CMTS, it again changes the Ethernet headers from Router MAC:CMTS MAC to CMTS MAC:Subscriber MAC .

This model shares many of the same values as the L2 Model. In addition, this model reduces the size of the router's ARP and ND tables compared to the L2 model, enhances the scalability due to the use of subnets, and provides an easier transition path for traditional CMTSs. A comparison of the two models is included in **Table 2** .

Table 2: Comparison of L2 and L2/L3 Models

Values	L2	L2/L3
Load balancing, failover, traffic control, premium services	X	X
Reduces need for routing protocols on each CMTS	X	X
No network renumbering during node splits	X	X
Reduced MAC learning at Routers (Hybrid OF/non-OF routers)		X
Enhanced scalability due to subnets		X
Easier transition for existing CMTSs		X

Implications for Cable

Both the L2 and L2/L3 models are viable for the CMTS. Regardless of the model, OpenFlow CMTSs offer benefits such as load balancing, failover, traffic management, and premium service support. Also, both models are capable of addressing the use cases described above. However, because of the simpler transition path for existing CMTSs, I recommend use of the L2/L3 model for initially phasing in OpenFlow support.

CONCLUSION

As we have discussed, OpenFlow is one piece of the overall SDN puzzle. As the most fully-developed SDN technology, and as the one with the most exposure, it is important to identify its place in cable networks.

The use case analysis described above identified several benefits OpenFlow brings to cable. First, it provides incremental enhancements to existing services such as L2VPN and lawful intercept. Second, as a step towards virtualization, it provides traffic control features and management tools. These enhancements were evident in the managed firewall, Carrier Grade NAT, IDS/IPS, and caching use cases.

Adding OpenFlow to an existing CMTS requires a mechanism to differentiate OpenFlow from non-OpenFlow traffic and a forwarding model. The hybrid service flow approach for traffic differentiation described above best fits the use cases discussed in this paper. Traffic on a specially marked service flow, or from a defined source MAC/IP address is processed by OpenFlow, while other traffic is forwarded using traditional methods. Also, the L2/L3 forwarding model provides a way to introduce OpenFlow to existing networks without disrupting established services.

REFERENCES

CableLabs BSoD L2VPN Specification, CM-SP-L2VPN-I10-121004, October 4, 2012, Cable Television Laboratories, Inc.

DOCSIS Cable Broadband Intercept Specification CM-SP-CBI2.0-I04-110224, February 24, 2011, Cable Television Laboratories, Inc.

ITU-T Recommendation Y.1731 (07/11), OAM functions and mechanisms for Ethernet based networks.

MEF Technical Specification 35: Service OAM Performance Monitoring Implementation Agreement, April 2012.

OpenFlow: Enabling Innovation in Campus Networks. Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, Jonathan Turner. March, 2008

Making Rational HFC Upstream Migration Decisions in the Midst of Chaos

Dean Stoneback, Fred Slowik
ARRIS

Abstract

It is estimated that approximately 1.5M miles of cable plant is now deployed in the United States. The majority of existing HFC networks have active 5-42 MHz reverse path in place. Much discussion has occurred during the past year relative to expanding reverse path network capacity on existing networks and how that expansion might proceed. Various ideas on the subject have included increasing the reverse pass band to 85 MHz, to 200 MHz, or even higher; adding a tri-split filter configuration with a second reverse path above 1GHz; relying upon node splitting and fiber deep migration to accomplish the goal; and incorporating a combination of solutions. Each approach has both technical and financial merits as well as disadvantages which need serious consideration.

To evaluate the costs of extending the upstream to 85 MHz or to 200 MHz (or greater) we will look at the issues involved in operating an HFC network upstream to those frequencies, including RF gain requirements, slope and AGC requirements and required laser performance. The cost of each of these components increases as the upstream bandwidth increases.

INTRODUCTION

Cable networks have been deployed universally with active reverse path since the inception of HFC in the 1980's. Essentially, the reverse coaxial path has been designed and operated for 5-42 MHz in North American networks and up to 5-65 MHz in some international regions. In HFC systems, due to the relatively small coaxial portion of the network, performance parameters such as CNR and NPR are dominated by the optical network.

Coaxial plant in HFC generally uses up to three basic types of low cost amplifiers. These are the four port, high output level amplifier (Type A), the intermediate multi-port amplifier (also called a mini-bridger – Type B), and the single port, low cost line extender (Type C). The same single gain hybrid reverse amplifiers are used in each type of product; however, the amplifier station's operating gain is influenced by the degree of internal loss devices such as diplex filters and port combiners required. These losses plus equalizer losses and perhaps optional thermal circuitry for controlling minor level variations due to temperature changes result in different station operational gains for each type of amplifier used. A Type A amplifier may have 17.5 dB, a Type B amplifier may have 20 dB and a Type C amplifier may have 24 dB of station operating gain with all losses considered.

The coaxial network is a mini-tree/branch network where individual reverse signal paths funnel into common signal paths returning to the node. For this reason, each amplifier contains reverse path equalization and attenuation capabilities located on the output side of the reverse amplifier in order to be able to properly align and balance the reverse network having signals returning from different originations. A similar situation exists from the subscriber's customer premise equipment (CPE). Each home can have several devices such as cable modems and set top boxes. These reverse CPE signals all enter the HFC network through a series of cascaded passive tap devices. Operating levels for these terminal devices are remotely controlled by the network's CMTS or other addressable control devices at the system's headend or hub location. CPEs do have maximum transmit power limitations which must be adhered to during the system design process. Issues such as excessive passive losses must be closely monitored in order to enable the reverse path to operate properly.

Today's existing networks with frequency limitations of 42 or 65 MHz generally have sufficient amplifier gain and CPE output power to insure proper reverse path operation provided the network was properly designed. In fact, it is estimated that 95% of reverse amplifiers contain attenuators due to an overabundance of reverse gain. In brownfield plant, such as is the majority of North American networks, and many international networks, upgrading the reverse path frequency limit needs to be reviewed in order to ensure proper products exist that are able to maintain the existing amplifier locations in the design. The last thing brownfield operators want is to have to re-plumb their networks. Also, the degree of change to the existing amplifier components may be impacted due to the reverse path upper frequency limit desired.

UPSTREAM LEVELS

Overview

Before beginning to select proper signal levels, it is essential to first understand how the return path works. Figure 1 shows a typical HFC Network (note that only the return path components are shown in the headend and fiber node). Signals originate in the home (1) and flow through the plant towards the headend. The signal level in the plant is determined by the RF level produced by the transmitter at the house, which is most often a cable modem.

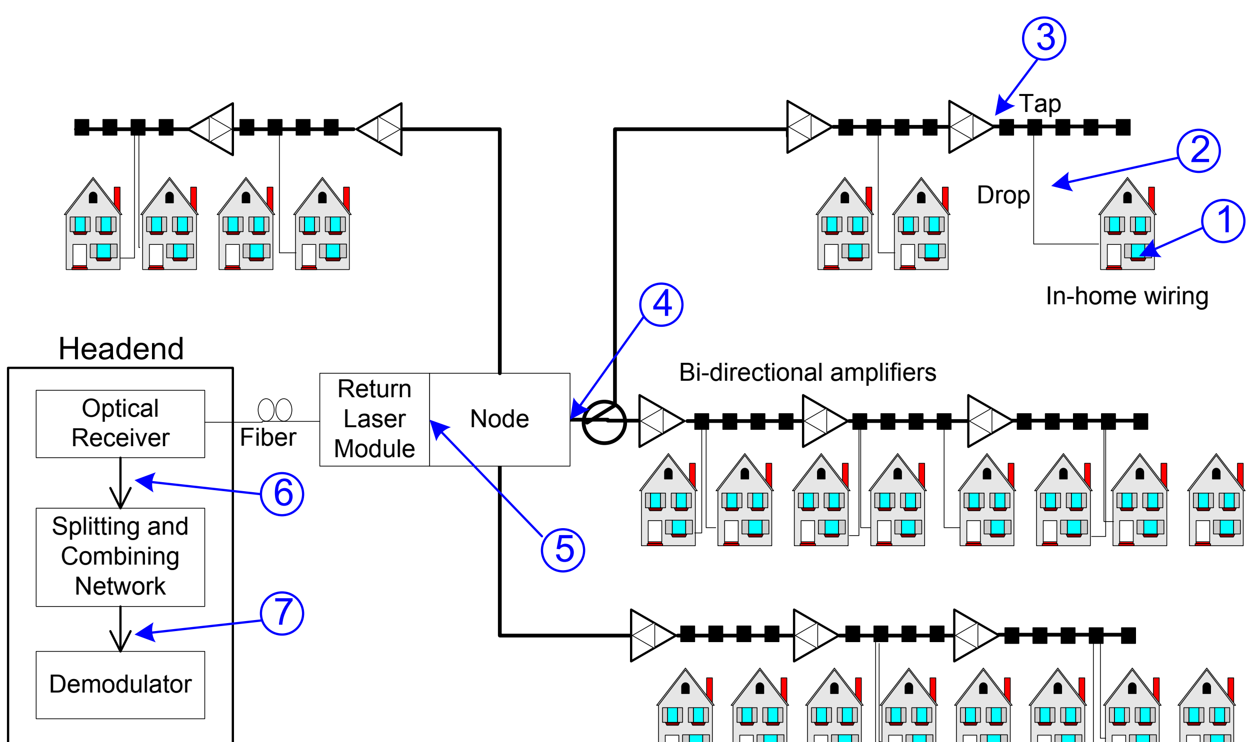

Figure 1: Upstream Levels in an HFC Network

After the signal leaves the cable modem, it goes through a variety of losses such as in-house cable, splitters, ground block, drop cable, tap port, and feeder cable before reaching the amplifier station port (3). All signals from the homes go through different amounts of loss, but all of these signals should arrive at the amplifier port (3) at the same level. This is a key premise of return path design. These variations require setting the transmitters in each cable modem to its own unique level – the level that produces the desired signal level at the amplifier.

Once the signals reach the amplifier, they continue on their way toward the headend. Every span of cable between two amplifier stations must be aligned to unity gain so that the return path gain of every amplifier station exactly matches the loss of the cable and passives following it (i.e., the cable span towards the headend). When the spans are all set to unity gain, the signal levels will be the same at every station. Ultimately, the signals reach the node station (4). Because the amplifiers have been aligned for unity gain, the signals at the node station port (4) are the same level as the signals at each amplifier station port (3). From the node station port, the signals continue on to the return path laser module (5). The relative levels between the node station port and the input to the laser are adjusted by selecting the proper gain or attenuation level in the node (see Figure 2).

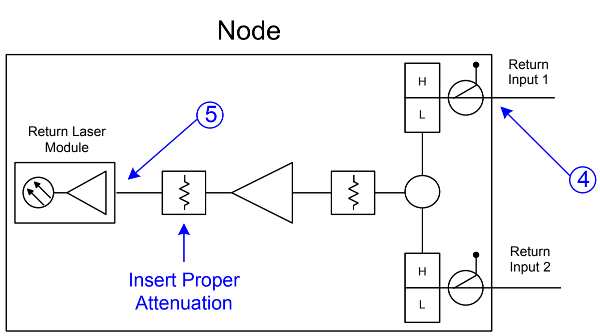

Figure 2: Adjusting Upstream Gain in the Node

After entering the return path laser module, the signals are then carried on fiber to the headend or hub location where they are converted back to RF by a fiberoptic receiver. This RF signal is then fed to the demodulator for that particular service. For DOCSIS services, the demodulator is the CMTS upstream input port.

Long Loop AGC

Many field technicians are already well-acquainted with aligning plants to unity gain, as described above. What is not immediately obvious, however, is how the real signals behave in a functioning plant. Long loop AGC refers to the process of adjusting the home signal levels via instructions from the demodulator in the headend. Most demodulators measure the level of the RF signal arriving at its input port (7). If this level is incorrect, a command is sent out via the forward path to the box in the house telling it to raise or lower its level accordingly. By making these adjustments, the demodulator assures that all signals from the plant arrive at the demodulator at the same level, within some accuracy limit.

"Long loop" refers to the fact that the commands are issued all the way from one end of the plant (the headend) to the other end of the plant (inside the home) to affect a signal originating inside the home and destined for the headend. Thus, a long loop is formed all the way across the plant and then back again. "Gain control" refers to the process of automatically adjusting gain. Strictly speaking, it is the level and not the gain that is controlled, but the term "long loop AGC" is widely used anyway.

Determining Ideal Upstream Levels

The following three different levels must be considered in every return plant:

- The level at the return path input port of every amplifier station (3) and (4)
- The level at the input to the return path laser module (5)
- The level at the input to the demodulator (7)

These three levels are based on independent criteria. The first step is to select the ideal level at the amplifiers based on the available transmitter power from the box in the home and the maximum loss through which that signal must travel on its way to the amplifier. Then, choose the ideal level at the return path laser module based on laser dynamic range and clipping versus carrier-to-noise (C/N) performance. Select the ideal level at the demodulator according to the manufacturer's specification. Once chosen, add gain or loss between each component so that all three simultaneously occur at the ideal level.

Levels in the plant are usually designed on a per-channel basis. The design begins with determining the available transmit power from the cable modem. Then the total loss between the cable modem and the amplifier station port is calculated. The transmit power minus the total loss is the available level at the amplifier station port. In most cases, some margin is added to that level to allow for loss variations over time and temperature.

The ideal level at the return path laser module (5) is a function of the type of laser and the module design. Several methods are available for determining the optimal input level to the module. In most cases, the manufacturer provides an optimal level. Generally, a trade-off exists between C/N and clipping distortion, which can best be identified by performing a Noise Power Ratio (NPR) test. In most cases, the ideal level determined for the laser is represented as composite total power.

Since the plant levels are generally calculated on a per-channel basis and the transmitter levels are calculated on a total power basis, a conversion must be performed. In most cases, operators desire that the upstream channels have a constant power spectral density, sometime referred to as constant power per Hz. Assuming all channels have the same bandwidth, one can convert from power per channel to total power by simply adding 10*log(number of channels).

Effect of Return Bandwidth on Upstream Levels

Since the total power at the laser transmitter is a function of power per channel and the number of channels, the total power will increase as the number of channels increases. Thus, wider return bandwidths, such as 85, 200 or 300 MHz will have larger total powers at the laser transmitter for the same individual channel levels.

Example of Upstream Levels

The first component to discuss is the output level available from cable modems. The DOCSIS spec for output level[i] is shown in Table 1. If the cable modem is only transmitting one upstream channel at a time, then it is capable of producing at least 57 dBmV. However, if it is a DOCSIS 3.0 modem and is transmitting multiple upstream channels, then guaranteed available maximum level per channel is reduced.

Number of Transmit Channels	Pmax (dBmV) TDMA
1	57
2	54
3-4	51

Table 1: Required Maximum Transmit Power for CMs for 64-QAM

The next component to analyze is the loss between the cable modem and the amplifier station port. Finally, the per-channel levels must be converted to a total power in order to properly calculate the required upstream node gain between the node station port and the upstream laser transmitter. An example is shown in Table 2.

	Levels	Power Increase Above 5-42 Power
CM Tx Level	52 dBmV	
Loss from CM to Tap Port	9 dB	
Largest Tap Value	23 dB	
Input to Amplifier and/or Node (per channel)	20 dBmV	
Total Power for 6 Channels (5-42 MHz)	27.8 dBmV	
Total Power for 12 Channels (5-85 MHz)	30.8 dBmV	3.0 dB
Total Power for 5-200 MHz	35.1 dBmV	7.3 dB
Total Power for 5-300 MHz	36.9 dBmV	9.1 dB

Table 2: Example of Upstream Levels and Conversion to Total Power

Table 2 assumes that the cable modem is able to transmit at 52 dBmV with some margin. Thus, this design does not allow for a DOCSIS 3.0 modem that is transmitting more than 2 channels at a time (see Table 1). Notice that the total power at the amplifier and node station ports will be 28 dBmV for a 5-42 MHz return bandwidth. This is a very common HFC design level. If the return bandwidth is increased to 85 MHz, the number of channels increases to 12 and the total power increases 3 dB to 31 dBmV. Similarly, if the return bandwidth increases to 200 MHz, the total power increases to 35 dBmV, which is 7.3 dB higher than it was for the 5-42 MHz return.

SELECTING THE OPTIMAL RETURN BANDWIDTH

Previous papers[ii,iii] have included a detailed analysis of the options available for increasing upstream bandwidth. This paper will focus on the costs of upgrading from a 42 or 65 MHz return to a 85, 200 or 300 MHz return.

85 MHz Mid-Split

85 MHz was selected years ago as the next likely maximum frequency for return path operation. The primary reason that 85 MHz was selected was so that the FM radio band, which operates from 88 to 108 MHz, would not be inside the return path operating band. Putting the FM band in the cross-over region reduces the likelihood that ingress from FM radio stations will be an issue.

Changing from 42 MHz or 65 MHz to mid-split is conceptually simple. All one needs to do is change the diplex filters in the nodes and amplifiers and realign the plant. In most cases, the amplifiers will have sufficient bandwidth and gain and the plant will not require any type of upstream AGC. When doing an upgrade, don't forget to change out any feederline equalizers that have diplex filters in them.

200 MHz High Split

Many people think that if 5-85 MHz is a good choice, then 5-200 MHz must be even better. However 5-200 MHz has some significant disadvantages.

Several of the disadvantages have nothing to do with the HFC network.

- A 5-200 MHz split means that downstream signals do not start until at least 250 MHz. Thus, a huge amount of deployed CPE cannot function and must be replaced. In particular millions of settop boxes have a downstream out-of-band receiver that cannot be tuned above 130 MHz.
- Loss of multiple VHF channels, some of which need to be carried on-channel.
- Reduction of downstream bandwidth

Other disadvantages related to the HFC network are:

- The entire FM band will now be in the return band. Thus, there is a the potential for large ingress.
- Return signals will exist in the aeronautical band. Leakage of return path signals becomes a concern.
- The change in gain of the coaxial network over temperature is no longer trivial.

Table 3 illustrates the gain and tilt change of 1000 feet of QR540 cable vs. temperature and frequency. One can see that the tilt from 5-42 MHz is less than 3 dB and that the change in gain and tilt across the full outdoor temperature range is less than 1 dB. Thus, amplifiers operating with a maximum upstream frequency of 42 MHz do not need gain and tilt correction to compensate for changes in temperature.

Conversely, the numbers for 5-200 MHz operation are not as optimistic. The tilt is almost 8 dB. This means that cable modems will need to transmit 8 dB higher if transmitting at 200 MHz than they would if transmitting at 5 MHz. This additional gain variance will be difficult to accommodate in the return path design. Perhaps even worse, the change in gain and tilt across the full temperature range is approaching 2 dB. Thus it is likely that some type of gain and tilt control will be required for 5-200 MHz operation through a cascade of amplifiers.

Loss of 1000’ of QR540 Cable from 5 to 300 MHz
(Loss at 750 MHz = 18 dB, Loss at 1002 MHz = 21 dB)

	25C		-40C		+60C		-40 to +60 Change	
Frequency (MHz)	Loss (dB)	Tilt (dB)	Loss (dB)	Tilt (dB)	Loss (dB)	Tilt (dB)	Loss (dB)	Tilt (dB)
5	1.49		1.68		1.38		0.30	
42	4.31	2.82	4.87	3.19	4.01	2.62	0.86	0.56
65	5.36	3.87	6.06	4.38	4.99	3.60	1.07	0.77
85	6.13	4.64	6.93	5.25	5.70	4.32	1.23	0.93
200	9.40	7.92	10.63	8.95	8.75	7.36	1.88	1.58
300	11.10	9.61	12.54	10.86	10.26	8.88	2.28	1.98

Table 3: Comparison of Cable Gain and Tilt Change vs. Temperature and Frequency

Forward-Driven Return AGC

One potential solution to allow 200 MHz upstream RF cascades is to drive a return path bode equalizer with the downstream bode equalizer control signal, as shown in Figure 3.

Although the control for such a system is relatively simple, there is minimal, if any, equipment on the market today with this functionality.

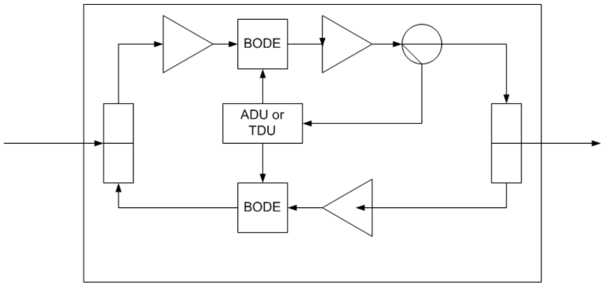

Figure 3: Forward-Driven Return AGC

Optical Link Performance

The most critical component in the upstream path in the HFC network is the laser transmitter. Receiver gain, output level capability, noise performance and distortion are also very important. All these components will have degraded performance when the upstream bandwidth is increased to 200 MHz. For instance, the noise power ratio (NPR) of an optical link for various bandwidths is shown in Figure 4 and summarized in Table 4.

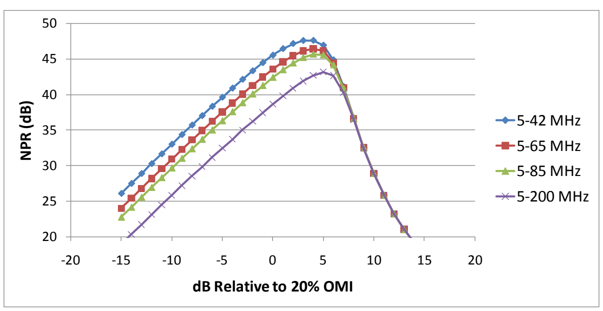

Figure 4: NPR of Upstream Link vs. Split

Bandwidth Increase	NPR (SNR) Reduction
5-42 to 5-65	2.10 dB
5-42 to 5-85	3.35 dB
5-42 to 5-200	7.22 dB

Table 4: NPR of Upstream Link vs. Split

ANALYSIS EXAMPLE

Theoretical Analysis

The reverse path network is comprised of a combination of cable loss and passive (flat) loss. Cable attenuation varies in accordance with the square root of the ratio of two frequencies. Passive loss does not follow that rule and is considered relatively flat across the reverse band frequencies under consideration in this paper. This means that the ratio of cable loss versus flat loss in the coaxial network impacts the amount of gain required in the reverse path. Table 5 illustrates a theoretical example of the reverse amplifier gain required based upon various combinations of cable versus flat loss – extending from 100% cable loss to 65% cable / 35% passive loss to 50% cable / 50% passive loss. It becomes quite evident that the more loss consumed by passive devices impacts the amount of reverse gain required in the amplifier. The bottom of Table 5 shows the total return path gain of typical amplifier stations. The gains are highlighted. The scenarios in the top portion of Table 5 for which that product has enough gain are also highlighted.

Table 5 indicates that under present conditions we would scarcely have enough reverse gain to accommodate an 85 MHz upgrade in applications where passive losses are present. As we all know, theory versus practice are often different and can produce two different outcomes. Theory does not factor in anomalies that exist in the real world such as short spaced amplifiers for placement optimization or the fact that most forward path amplifiers contain input pads (5dB on average) for example. Hence, we also reviewed three practical design applications to validate these data.

			Forward			Reverse				
Frequency (MHz)			**1000**	**870**	**750**	**300**	**200**	**85**	**42**	**5**
P3-625 loss/ft	**@68deg. F.**		0.0207	0.0193	0.0179	0.0113	0.0093	0.0060	0.0042	0.0015
Station Gain - dB (A/B)	**100%**	**Cable**	42	39.2	36.4	23.0	18.8	12.2	8.6	3.0
Station Gain - dB (C)	**100%**	**Cable**	34	31.7	29.4	18.6	15.2	9.9	7.0	2.4
Station Gain - dB (A/B)	**65%**	**Cable vs. FL**	42	40.2	38.3	29.7	26.9	22.7	20.3	16.6
Station Gain - dB (C)	**65%**	**Cable vs. FL**	34	32.5	31.0	24.0	21.8	18.3	16.4	13.5
Station Gain - dB(A/B)	**50%**	**Cable vs. FL**	42	40.6	39.2	32.5	30.4	27.1	25.3	22.5
Station Gain - dB (C)	**50%**	**Cable vs. FL**	34	32.9	31.7	26.3	24.6	22.0	20.5	18.2
Current Station Gains	**A**		**42**						**17.5**	
(Include 3dB average fwd. pad.)	**B**		**42**						**20**	
	C		**34**						**24**	

Table 5: Calculation of Required Return Amplifier Gain

Practical Analysis

The purpose of this section of this paper is to present some of the issues that arise at different reverse path frequencies in real world applications as well as to understand the cost implications.

Reverse path amplifiers generally have five basic elements that are impacted by a reverse path frequency upgrade. These are station gain, diplex filters, attenuators, equalizers, and thermal control. As previously stated, amplifiers used in current 42 and 65 MHz reverse plant have more than sufficient gain to operate properly. Some of those amplifier hybrids now in place are capable of 5-200 MHz operation so if they have sufficient gain at these upgrade frequencies, theoretically they may not need replacement. Other components such as diplex filters, attenuators, equalizers and thermal control units are frequency dependent and would need to be replaced however. This leads to the analysis of identifying which elements need replacement depending upon desired reverse upgrade frequencies. Also, since this upgrade requires a truck roll, it may be important to consider what the desired end game is for reverse frequency versus the perceived life of the network. In an effort to answer these questions, we reviewed several network designs in order to determine the impact of upgrading existing 5-42 MHz reverse path to 5-85, 5-200 and 5-300 MHz reverse.

Prior to beginning this analysis, we first need to establish the design parameters followed in the reverse path designs at 5-42MHz. The output level available from cable modems was shown in Table 1. A complete list of design parameters is presented in Table 6.

Parameter	Common Existing Reverse Specifications	Comments
Cascade:	Up to N+5	
Amplifier Types:	Types A, B and C	

Parameter	Common Existing Reverse Specifications	Comments
Node Reverse Input:	Up to 20 dBmV/channel	This was a common level for DOCSIS 1 & 2 era designs. DOCSIS 3 presented a reduction in CPE output of 4 dB so this level was reduced by 4 dBmV/channel.
Amplifier Reverse Input:	Up to 20 dBmV/channel	Same as above.
Largest Tap Value Deployed:	23	Conditioning taps (with equalizers or cable simulators) may be used to correct excessive positive or negative slope.
Reverse Tap Port Minimum Input Level:	46 dBmV/channel	Level reduced to 42 dBmV in DOCSIS 3 designs.
Reverse Path Drop Loss CPE to Tap Port:	9dB	Combination of drop cable and passive loss.
CPE Maximum Output Per Channel:	55 dBmV/channel	Level reduced to 51 dBmV for DOCSIS 3 CPEs.

Table 6: Upstream Signal Level Assumptions for 5-42 MHz

Due to the reduction in CPE output level of DOCSIS 3.0 CPEs we began the exercise by adjusting the original design 5-42 MHz parameters as illustrated in the Comments section of Table 6 for the three sample designs. This resulted in negligible impact to the reverse designs.

DESIGN SCENARIOS

Three different sample design areas were reviewed along with three reverse design upgrade scenarios having an estimated cable to passive loss ratio average of 65%/35%.

Scenario 1, increased reverse bandwidth from 5-42 MHz to 5-85 MHz using the adjusted 5-42 MHz DOCSIS 3.0 design parameters. This resulted in minimal change to the reverse design.

- The amplifier to amplifier gains were sufficient to hold locations using existing hybrids.
- Amplifiers required change out of reverse equalizers, attenuators and diplex-filters.
- CPE to amplifier/node inputs held but approximately 12% of tap face plates needed to be changed. The alternative would be to either increase CPE output by 1 dB or reduce amplifier/node inputs by up to 1 dB. The latter would result in a reduction of reverse CNR and NPR.
- Additionally, if reduced input levels are under consideration, it is important to verify that the existing nodes in place contain sufficient reverse optical transmitter gain needed to drive reverse path optical links.
- Reverse thermal compensation modules continued to be used as the level variation over temperature was minimal.

Scenario 2 increased reverse bandwidth to 200 MHz. In this case, network revisions became more extensive.

- The amplifier to amplifier gains were sufficient to hold locations using existing amplifier stations.
- Amplifiers did require change out if reverse hybrids were not 200 MHz capable. Reverse equalizers, attenuators

and diplex filters also needed to be changed.

- CPE to amplifier/node inputs did not hold and up to 90% of tap face plates needed to be changed resulting in significant network redesign (and possible amplifier respacing). The alternative would be to either increase CPE output by 4 dBmV or reduce amplifier/node inputs by up to 4 dB. At this amount of input reduction, CNR and NPR performance in the reverse path can be reduced and the plant levels are getting closer to the level of ingress noise.
- Additionally, if reduced input levels are under consideration, it is important to verify that the existing nodes in place contain sufficient reverse optical transmitter gain needed to drive reverse path optical links. As shown in Table 2, the total power will increase as the bandwidth increases, thus less return gain will be required.
- Reverse thermal compensation modules would likely need replacement as well to accommodate increased level variations over temperature as Table 3 illustrated.

Scenario 3 increased reverse bandwidth to 300 MHz. In this case, network revisions became even more extensive.

- The amplifier to amplifier gains fell short by up to 5 dB using existing amplifier stations.
- Amplifiers did need to change out as well as reverse equalizers, attenuators and diplex filters.
- CPE to amplifier/node inputs did not hold and up to 90% of tap face plates needed to be changed resulting in severe network redesign. The alternative would be to either increase CPE output by 5 dBmV or reduce amplifier/node inputs by up to 5 dB. At this amount of input reduction, CNR and NPR performance in the reverse path can be reduced and the plant levels are getting closer to the level of ingress noise.
- Additionally, if reduced input levels are being considered, it is important to verify that the existing nodes in place contain sufficient reverse optical transmitter gain needed to drive reverse path optical links. As shown in Table 2, the total power will increase as the bandwidth increases, thus less return gain will be required.
- Reverse thermal compensation modules would likely need replacement possibly with automatic gain control depending upon the cascades in the network as Table 3 illustrated.

Table 7 summarizes these results a bit more concisely. The net result is that although a 5-85 MHz reverse upgrade can be achieved with minimal impact, moving to 5-200 MHz or 5-300 MHz results in far more network modification required.

Item	Scenario 1 85 MHz	Scenario 2 200 MHz	Scenario 3 300 MHz	Comments
Amp to Amp Gain	Held	Held	5 dB low	S2 may require amp replacement. S3 will require amplifier replacement.
CPE to Amp Inputs	Held	4 dB low	5 dB low	Requires increased output CPEs for S2 & 3.

Item	Scenario 1 85 MHz	Scenario 2 200 MHz	Scenario 3 300 MHz	Comments
% Tap Faceplate Change	12%	Up to 90%	Up to 90%	S2 & 3 cause significant rework and likely additional amplifiers. Increased CPE outputs would alleviate this.
Reverse Amp Reuse	Yes	Possibly if 200 MHz capable.	No	
Pad & EQ & Diplex Change	Yes	Yes	Yes	
Temperature Control	Thermal	Possibly AGC	AGC	
Truck Roll Required	Yes	Yes	Yes	

Table 7: Comparison of Upgrade Scenarios

Table 8 reviews the constructed cost implications of each scenario. It contains several assumptions made. Cost is based upon tap faceplate change out as opposed to replacing CPEs with higher output devices. The "Amp Accessories" line includes diplex filters, pads and EQs. In the 85 MHz case, the gain stages were assumed to be reusable. In the 200 and 300 MHz cases, we assumed the gain stages needed to be upgraded.

We also observed that 5-200MHz and 5-300 MHz reverse upgrades with tap faceplate change out would likely cause significant network re-plumbing and create the need for new active device locations. This could then violate the network powering structure.

Item	Scenario 1 85 MHz	Scenario 2 200 MHz	Scenario 3 300 MHz
Total Reverse Upgrade Price/Mile	$1203	$4598	$4598
Replace Reverse Hybrid	$0	$240	$240
Add New Amplifier Location	$0	$600	$600
Amp Accessories	$264	$385	$385
Tap Face Plates	$39	$293	$293
New P.S.	$0	$80	$80
Labor	$900	$3000	$3000
New CPE	$0	$0	$0

Table 8: Cost Comparison of Upgrade Scenarios

TRI-SPLIT 1200MHZ REVERSE

The results portrayed in the 200 MHz and 300 MHz reverse upgrades demonstrated that higher gains and AGC are required to operate with high-bandwidth splits. It is evident that attempting to move to a tri-split filter with 5-42 MHz and 1100-1200 MHz return would be even more costly. Forward amplifier gain would experience additional loss due to triplex filter loss. Reverse amplifier gains would need to exceed the current 42dB maximum forward gains now deployed at 1000 MHz. Due to the increased reverse output level requirements, issues with crosstalk would be likely; therefore, a complete new e-pack and perhaps amplifier housing (depending upon existing housing capabilities) would be required. Cost of this model is deemed excessive.

AMPLIFIER UPGRADE METHODS

Most HFC networks deployed today use a 5-42 MHz or 5-65 MHz return. However, most operators are seriously considering moving to a 5-85 MHz or higher upstream bandwidth network in the near future. Operators want to deploy a product today that can serve their needs in the future. To do this, there are several options:

- Have multiple diplex frequencies in the initial product with some type of switching mechanism to select the desired frequency. The goal is to affect future change without visiting the amplifier.
- Have a pluggable sub-module that can be replaced in the future. The goal is to do the upgrade quickly without needing to discard or bench-upgrade existing product.
- Perform an electronics package (EPAC) swap in the future.

The following sections will consider each of these options.

Multiple Diplex Frequencies in Initial Product

Having multiple diplex frequencies in the initial product with some type of switching mechanism to select the desired frequency has several advantages, including:

- No need to visit amplifiers in the future
- Low down time
- Goal of no part changes or craft issues
- No bench upgrade

Unfortunately, there are several significant disadvantages with this approach:

- The “final” frequency is not known today. In particular, the selection of exactly “200 MHz” has not yet been decided by the industry.
- Highest initial product cost
- Increased product complexity
- Requires sweep and balance at 200 MHz during initial installation, which requires vacating all DS frequencies up to at least 200 MHz, then a revert back to 85, 65 or 42 MHz
- Risk that it “won’t work” when switched years from now
- Significant HFC plant changes may be required day 1 to accommodate highest upstream split increment incorporated

Pluggable Sub-Module

- Having a pluggable sub-module that can be replaced in the future, with a goal of doing the upgrade quickly without needing to discard or bench-upgrade existing product has several advantages, including:
- Allows plant bandwidth changes to occur when needed
- Low down time. Pre-configured modules plug in quickly.

Some disadvantages of this approach are:

- Need to visit the amplifier
- Higher initial cost with future incremental cost

- More challenging / higher risk design effort
- May still require powering down the feeder leg during the upgrade
- Will require sweep and alignment adjustments to the host EPAC
- Host EPAC design and maximum frequencies are locked down on day 1

EPAC Swap

Performing an electronics package (EPAC) swap in the future has many advantages including:

- Lowest initial cost
- Pay as you grow
- Lowest product complexity
- Option to either reconfigure or replace EPACs during the upgrade, depending on product age
- Don't need to make a bet today on the future configuration
- Very low down time. Pre-configured and tested EPAC modules are plugged in

There are a couple disadvantages to this method:

- Need to visit each amplifier during upgrade
- Could require bench top configuration and alignment and then bicycling of existing EPACs

The EPAC swap appears to be the best method for operators to upgrade their networks in the future. The other methods are more expensive, more intrusive during initial setup and include significant risk that the decisions made during initial deployment will not be the correct configurations in the future.

Plant Upgrade Procedure

When it is time to upgrade from one frequency split to another split, the following procedures should be followed.

"Cold swapping" is the preferred method. To perform a cold swap:

- Power down and cold swap modules
- Power up and sweep and balance each node segment

If the system down time cannot be tolerated and an "in cascade" module upgrade is chosen, care must be taken to avoid any possible loop gain oscillation as follows:

- The network should be void of any RF sources from the upper end of the original return band-pass split to the lower end of the new downstream band pass
- The operational gain in both directions should be equalized and padded with the design values prior to module power up
- Extra care should be taken on short spaced amplifiers to make sure the gains are not too high
- Each node segment should be committed to and completed timely
- N split modules should not be intermixed in plant of other splits

CONCLUSIONS

Operators should plan for an upgrade to 85 MHz. 5-85 MHz reverse upgrades in properly designed networks require little network modification to accomplish. Return path bandwidths beyond 85 MHz significantly add to the expense of the network.

There is no clear market driver that indicates return bandwidths beyond 85 MHz will be required any time in the next 10 to 15 years. There is no clear standard on what the next incremental return path frequency beyond 85 MHz will be.

200 MHz and 300 MHz reverse upgrades increase cost and complexity throughout the plant. 5-200 MHz and 5-300 MHz reverse upgrades are more involved, requiring higher gain return path amplifiers with higher output power capacity and potentially requiring

active upstream gain and tilt control. Amplifiers without these features may need to be replaced. Significant tap face plate change out is also required unless CPE output levels are increased or input levels to the amplifiers and nodes are decreased, which will impact carrier-to-noise and carrier-to-ingress performance. Perhaps most importantly, there are tens of millions of deployed CPE units that require a downstream communication channel at frequencies lower than 130 MHz[iv], preventing an upgrade to a 200 or 300 MHz upstream without replacing the CPE equipment.

CPE output levels are a key contributor to the overall cost of reverse upgrades.

Operators want to deploy a product today that can serve their needs in the future. The best method to plan for a future upgrade is to deploy a product that can be upgraded in the future.

ACKNOWLEDGEMENTS

The authors would like to thank Mike Aviles and Stuart Eastman for their calculations and design evaluation contributions to this paper.

---

[i] CM-SP-PHYv3.0-I10-111117, “Data Over Cable Service Interface Specifications DOCSIS® 3.0 Physical Layer Specification,” available at http://www.cablelabs.com/cablemodem/specifications/specifications30.html

[ii] Rewriting the Book of Return – A New 10-yr Plan, Dean Stoneback, SCTE Cable-Tec Expo, October, 2010

[iii] Characterizing and Aligning the HFC Return Path for Successful DOCSIS 3.0 Rollouts, Robert Howald, Phillip Chang, Rob Thompson, Dean Stoneback, Vipul Rathod, Charles Moore, SCTE Cable-Tec Expo, October, 2009

[iv] ANSI/SCTE 55-1 and ANSI/SCTE 55-2, “Digital Broadband Delivery System: Out of Band Transport", available at http://www.scte.org/

Making Room for DOCSIS 3.1 and EPoC – Is your cable plant ready for an OFDM world?

Author - Phil Miguelez
ARRIS

Abstract

New standards are winding their way through CableLabs and IEEE that will eventually provide cable operators the ability to offer greatly increased data rate capacity to both residential and business customers. DOCSIS® 3.1 and EPoC will usher in a new modulation format for cable MSO's that will help to significantly close the current gap in digital capacity between FTTP and HFC service providers. A major challenge for the success of these next generation technologies is integrating new dedicated bandwidth segments into already constrained RF spectrum. This is particularly true for the upstream where the current 5 to 42 MHz channel allocation is already extremely limited.

Taking full advantage of the efficiencies related to OFDM transport without cannibalizing existing revenue generating RF spectrum will drive new requirements for expanded bandwidth optical and RF components. Although the initial deployment intent of DOCSIS® 3.1 is complete compatibility with the current 1 GHz RF bandwidth, many consider 1.2 GHz to be a logical end point that will maintain the full legacy HFC bandwidth as well as a new 200 MHz sub band for D3.1 or EPoC. Further expansion beyond 1.2 GHz is also a possible consideration for the future, allowing data rates up to 10 Gbps.

This paper will examine the impact of D3.1 and EPoC on current and future access plant components including headend lasers, nodes, and RF actives as well as taps and passives. Network design considerations, operating levels and system performance as the channel loading migrates to include OFDM will also be studied.

INTRODUCTION

Over the past several years' cable MSOs have very successfully launched voice and internet IP services across their HFC networks. DOCSIS QAM has been a large part of this broadband success story. But as competition from Telco and over-builders began to challenge the established cable markets, operators have felt the pressure to increase data rate capacity in order to meet the inevitable comparisons between HFC and fiber to the home (FTTH) networks. The DOCSIS cable standard has also continued to evolve from its early implementation to the current 3.0 standard, offering higher download and upload speeds. But the accelerating growth curve of IP data delivery still threatens to outpace the capacity of traditional DOCSIS transport. The well-publicized Nielsen data rate curves and CAGR plots continue to predict that cable operators will run out of bandwidth in a relatively short time unless some major system changes occur. Node segmentation and analog reclaim have provided breathing room for many operators, extending the available bandwidth to each subscriber. But new challenges to the dominance of cable broadband continue to surface coming from government initiatives, rapidly evolving technology, and the requirements of new business services customers.

THE NEED FOR SPEED

On March 16th 2010 the FCC published the National Broadband Plan[1]. The plan sets forth a number of goals targeting service improvements in both wireline and wireless access. Within the plan objectives are defined timelines to achieve specific download and upload data rate targets. For wireline residential access networks such as HFC cable and fiber to the home the first of these goals includes a minimum of 100 Mbps download and 50 Mbps upload speeds available at an affordable cost to at least 100 million homes in the US by the year 2020. Another goal specifies that every American community should have affordable access to at least 1 Gbps broadband service at institutions such as schools, hospitals, and government buildings. More recently, Julius Genachowski the chairman of the FCC issued a challenge to broadband providers calling for the deployment of gigabit Ethernet service in at least one community in each of the 50 states by 2015.[2]

To meet these growing challenges the IEEE 802.3 Ethernet working group issued a call for interest in November 2011 titled – "Operating the EPON Protocol over Coaxial Distribution Networks". Two months later, the "IEEE 802.3 EPON Protocol over Coax (EPoC) Study Group" was created. In June 2012, CableLabs the non-profit cable industry consortium, initiated a new specification effort to establish the requirements of next generation DOCSIS 3.1.

Downstream data usage rates have been growing at 50% compounded annual rates for several years. If this trend continues the subscriber data capacity needed within the next 10 years will exceed 10 Gbps.

The DOCSIS 3.1 specification will define a new modulation standard for HFC networks with a data rate capacity of 5 Gbps downstream (DS) and up to 1 Gbps upstream (US) while maintaining the current 1 GHz RF bandwidth capabilities of existing cable plant.

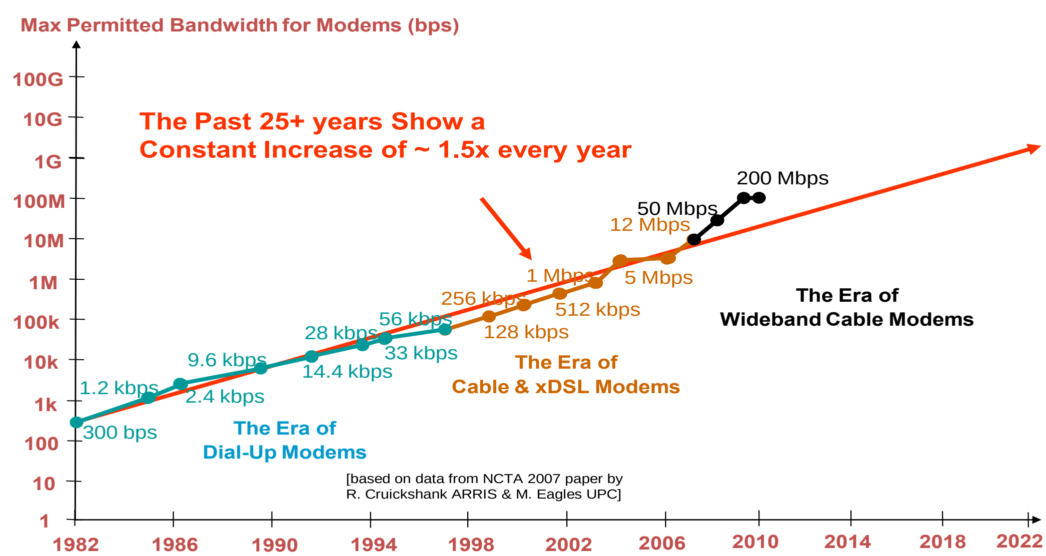

Figure 1 – Nielsen Curve for traffic growth over HFC Networks (3)

The potential for 10 Gbps downstream capacity is achievable with an RF spectrum expansion to approximately 1.5 or 1.7 GHz. Dramatic upstream capacity increases provided by DOCSIS 3.1 will also require significant RF bandwidth changes. High splits of 200 MHz to 400 MHz are still being debated to raise the upstream delivered data capacity to 1 Gbps or higher.

CableLabs has set a target goal to complete the D3.1 specification by year end 2013. Potentially this would allow modem chip sets to be developed and initially introduced as early as 2014 with CPE deployments following in 2015. The IEEE EPOC working group has also been meeting since the beginning of the year along with a number of ad hoc groups that are focused on specific PHY and MAC layer portions of the standard. The estimated timeline to complete the EPOC specification is currently late 2014 or early 2015. The large gap between specification delivery timelines of the two organizations is due to the differences in their respective charters. CableLabs is primarily accountable to its cable operator membership and only creates specifications for the MSO community. IEEE is an international standards organization that must obtain consensus across a wide range of users in many countries.

DOCSIS 3.1 and EPoC Development Goals

A primary goal of both the DOCSIS 3.1 and EPoC specification efforts is the capability to deliver spectrum efficient gigabit data rates. [(4)] To achieve this one of the first considerations is the selection of a modulation format. With a pre-existing transport network and limitations on the usable RF frequency bandwidth, the modulation format (channel width, modulation order, single carrier, multi-carrier, etc.) is the only dimension available to significantly increase the efficiency of the coaxial access link. Both working groups quickly focused on Orthogonal Frequency Division Multiplexing (OFDM) as the successor to single carrier QAM. OFDM is a multi-carrier format with each sub carrier modulated using higher order QAMs.

OFDM subcarrier modulation up to 4096 QAM allows 5 Gbps data rates using approximately 500 MHz of RF spectrum. The result is a 35% improvement in bit/Hz efficiency compared to DOCSIS 3.0 transport. The large bit/Hz efficiency increase could be used in place of node segmentation to improve bandwidth per subscriber - potentially at lower cost. Further expansion of data capacity to 10 Gbps will require an extension of the downstream bandwidth to at least 1.5 GHz and possibly higher.

A major goal for the D3.1 spec is the requirement to operate in existing HFC plant architectures. The common assumption is that downstream D3.1 channels will be placed at the upper end of the available frequency bandwidth above the existing broadcast and narrowcast channel lineups. Depending on the age and quality of the network this could include spectrum with higher frequency roll-off tilt, flatness variations, and degraded return loss performance. The spread spectrum nature of OFDM is more robust to these conditions due to the ability to adaptively modulate individual sub carriers.

Upstream goals for D3.1 include CMTS backward compatibility with D3.0 and D2.0 modems. It is also hoped that initial CCAP platforms that are just starting to be delivered will be able to be upgraded to D3.1 through a firmware revision or card change. Nearly 100% of North American cable networks currently use a 42 MHz return bandwidth. To achieve the full 1 Gbps data rate capacity of DOCSIS 3.1 an RF bandwidth of at least 200 MHz is essential. But as stated previously, a goal of D3.1 is that an upgrade is not a requirement for implementation. This maintains existing plant equipment use but limits the immediate impact of D3.1 in the upstream depending on the current US - DS frequency split. The DOCSIS 3.1 working group has also eliminated the idea of doing a top split where upstream spectrum would be placed above the downstream bandwidth.

While EPoC shares many of the same first order target goals as DOCSIS 3.1, a significant difference is that the IEEE standards group specification goal is to provide symmetric and asymmetric full duplex Ethernet transport over coax with no substantive changes to other EPON layers. In this case EPoC would only coexist with HFC. The EPoC transmissions would traverse between an EPON OLT and a coaxial network unit (CNU) modem at the subscriber termination side. In order to achieve symmetric data rates, EPoC transmissions could use either Frequency Domain Division (FDD) or Time Domain Division (TDD). FDD for symmetrical data rates of 2.5 Gbps or more would exceed the available RF spectrum of existing cable networks assuming there were no HFC channels carried on the same system. TDD would solve this problem for networks that plan to overlay HFC with EPoC at 1 Gbps data rates or higher.

Single Carrier vs. Multi-Carrier

DOCSIS 1.0 through the current DOCSIS 3.0 standards have all been based on single carrier QPSK or QAM formats. Single carrier modulation (SCM) uses a fixed, uniform modulation profile. The transmission performance is dependent on the signal to noise characteristics of the channel frequency. With defined channel bandwidths of 6 and 8 MHz, increasing the data capacity is achieved by increasing the QAM modulation order (8 bits/256 QAM, 10 bits/1024 QAM, etc.) of the transported channels or channel bonding. Increasing the modulation order requires an appropriate SNR level maintained across the entire channel(s) bandwidth. Other limitations of SCM are the complexity of bonding multiple channels and the performance degradation impact of noise and discrete interferers within the channel(s).

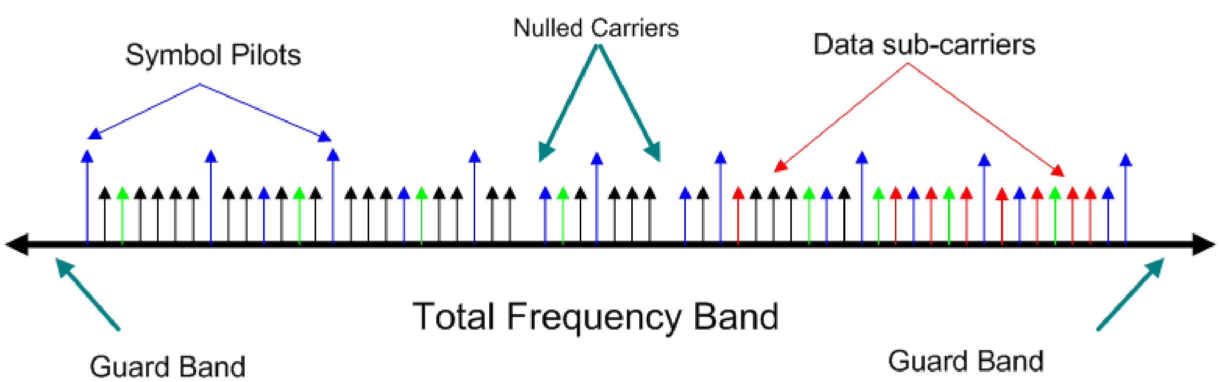

Figure 2 – Representation of OFDM Multi-Carrier Modulation

Multi-carrier modulation uses discrete multiple tones (DMT) spread across a wider frequency bandwidth that does not necessarily have to be a contiguous channel. Orthogonal Frequency Division Multiplexing (OFDM) uses multiple narrow 20 KHz to 50 KHz subcarriers that are each modulated with higher order QAMs. An OFDM FFT block of 192 MHz is the most commonly referenced channel width target for D3.1 and EPoC. DOCSIS 3.1 has further defined a minimum sub block channel size of 24 MHz. The 24 MHz minimum channel size was selected in order to have a common bandwidth allocation for DOCSIS and Euro DOCSIS. Since 24 MHz is a common denominator for both DOCSIS 6 MHz channels and Euro DOCSIS 8 MHz channels. Table 1 details the raw and estimated delivered data capacity for a 192 MHz FFT block and the various modulation formats that could be transported. An overhead efficiency factor of 30% was used in the data rate calculations below (Table 1). Estimates of this efficiency factor range anywhere from 20 to 35% depending on the source. The data capacity for other channel widths can be approximated as multiples of 24/192 MHz assuming the remaining sub-carriers have been nulled out.

OFDM DS Data Rate Capacity

QAM	Bits/symbol	FFT Block Sym rate (Msps)	Block size (MHz)	Raw Capacity (Mbps)	Efficiency (estimate)	Delivered Capacity (Mbps)
64	6	192	192	1152.00	0.7	806.40
256	8	192	192	1536.00	0.7	1075.20
1024	10	192	192	1920.00	0.7	1344.00
4096	12	192	192	2304.00	0.7	1612.80

Table 1 – Downstream Capacity Calculations for D3.1 OFDM

An advantage of OFDM is that the subcarrier modulation order can be varied across the channel to compensate for differences in SNR with frequency as shown in Figure 3. This feature allows OFDM to operate in links where the frequency gain response is not uniform due to passive losses or RF active performance. Individual subcarriers can also be nulled out in the case of discrete interfering signals. This allows OFDM to provide higher throughput than single carrier QAM under non ideal link conditions.

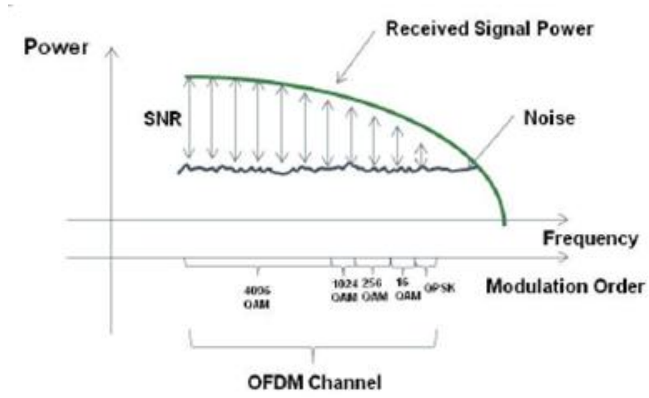

Figure 3 – Example of Adaptive Modulation Order for D3.1 OFDM

OFDM is less complex than MAC layer channel bonding, allowing easier scaling to higher data rates. OFDM is also more resilient to micro reflections, impulse noise, and ingress. In addition to these benefits, both the D3.1 and EPoC specification working groups are planning to also change the current forward error correction (FEC) scheme. Earlier versions of DOCSIS have all used Reed-Solomon FEC coding. In the early 1990's newer turbo FEC codes were developed that demonstrated improved efficiency in high noise channel environments. The discovery of turbo codes led researchers to look for other lower complexity coding solutions. These efforts resulted in the rediscovery of LDPC codes, first proposed by Robert Gallager(5) in his 1960 doctoral dissertation. Low Density Parity Check (LDPC) codes provide FEC solutions that are even closer to the Shannon capacity limit than any previous code. The improved spectral efficiency allows higher order QAM transmission at SNR levels that are 7 to 10 dB lower than achievable with traditional Reed-Solomon coding. Therefore, with LDPC the SNR needed to transport 1024 QAM is approximately equivalent to the DOCSIS 3.0 SNR for 256 QAM. Figure 4 shows a simulation of SNR values for different modulation orders and FEC levels.

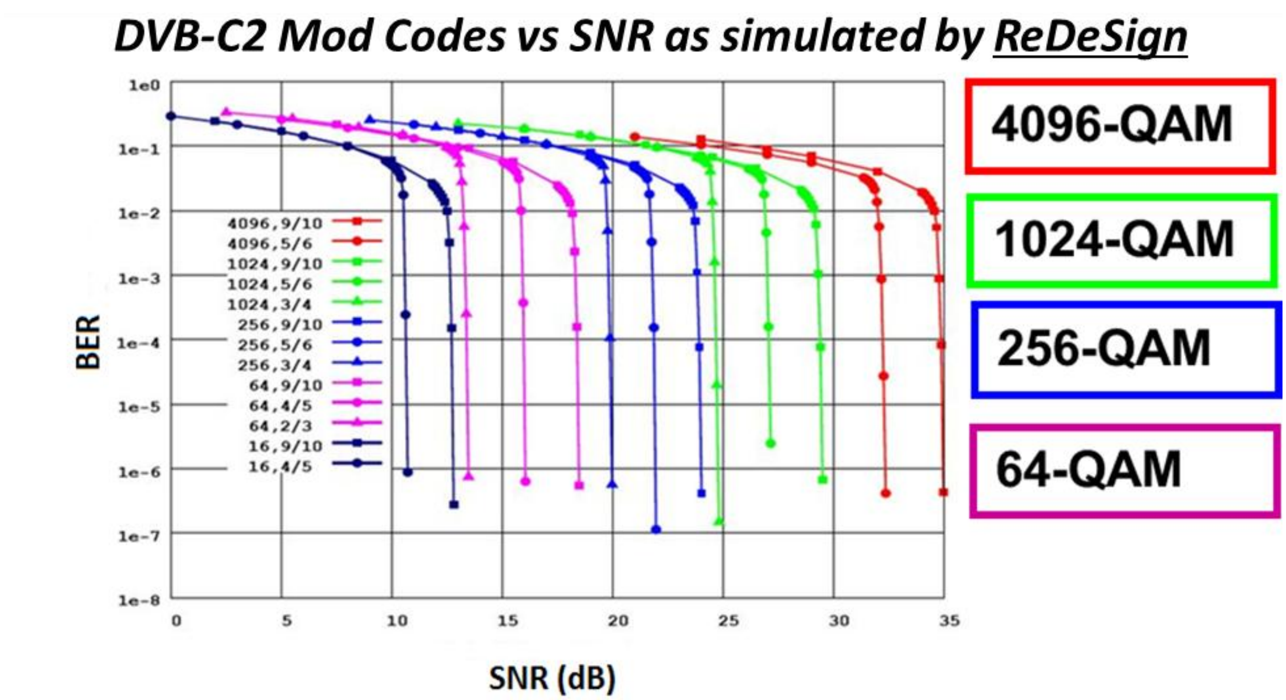

Figure 4 – SNR for OFDM with LDPC Forward Error Correction (6)

The efficiency of OFDM combined with LDPC FEC alleviates the need for higher CNR performance optics and RF actives when compared to single carrier 1024 QAM channels. A 1K OFDM channel should be able to maintain the same -6 dB from virtual level derate used today for downstream 256 QAM channels. This should allow systems that need to carry analog plus QAM channel loads to potentially add DOCSIS 3.1 or EPoC channels as long as the total RF input drive level to the transmitter is maintained. The SNR requirements for higher order modulation OFDM channels are still within the range of typical HFC networks but will have reduced margin against the normal range of network and seasonal variations. In this case the 4K OFDM channels could still be carried on legacy access links at the same -6 dB derate as existing QAM channels with the assumption that the modulation order and peak data rates would be backed off to areas of the network with lower SNR values. Changing the derate to -3 dB or higher as an example would buy back most of the lost margin for 4K OFDM transport but the increased power load could push the laser transmitter or RF amplifier into compression degrading the performance of the entire serving area link. More investigation is needed once the working groups have completed their spec definition efforts to determine the worst case loading conditions and the active device peak power performance that will be needed.

Discrete non time varying interferers could disrupt OFDM subcarriers that fall on the same frequency. In most cases these subcarriers can be nulled out with very little impact to the overall data rate of the OFDM block. In the case of mixed analog video channel loading with OFDM the concern is the number of CTB beats that will impact the D3.1 subcarriers. For a 192 MHz block of OFDM subcarriers, only 2% to 5% will be impacted by CTB beats generated from a 79 analog channel load. As the analog carriers are reclaimed this becomes less of a problem. A reduction from 79 analog channels to 60 channels results in a 6 dB reduction in CTB. Interleaver coding may also help reduce the impact of distortion beats such as CTB and CSO generated by analog carriers. Interleavers are typically used in multi-carrier wireless applications to mitigate selective signal fading by distributing the transmitted bit-stream across a wider range of frequencies rather than concentrating the bits on a narrow band of subcarriers. CTB beats are predictive based on the channel relationships of the analog carriers allowing interleaver algorithms to minimize the loss of critical parts of the bit-stream. The eventual transition to all digital loading by reclaiming the remaining analog video channels will completely eliminate the issue of CTB impairments.

OFDM in the Upstream

The legacy upstream bandwidth allocation is much more constricted than the downstream with less than 37 MHz available in most North American systems today. It is also anticipated that cable operators will maintain the current D2.0 and D3.0 channels for a considerable time, consuming a large portion of the limited clean spectrum in the 15 to 42 MHz bandwidth segment. The HFC upstream environment contains many more local sources of potential interference than the downstream. The SNR levels received from each subscribers' home has a wide distribution resulting from varying loss budgets depending on the tap position along the access coax path, ingress levels, and in-home wiring losses. To counter the dynamic nature of the upstream a variation of OFDM has been selected for subscriber premise equipment. Orthogonal Frequency Division Multiple Access (OFDMA) provides a combination of frequency domain and time domain multiple access by assigning different numbers of subcarriers to different users as shown in Figure 5 below. In addition to providing the same robustness to ingress and impulse noise as OFDM, OFDMA also enables adaptive modulation for every individual user. The modulation order can be dialed back to optimize throughput for subscribers with poor SNR values without affecting the upload speed of other customers on the same link.

The target FFT block size for upstream OFDMA is 96 MHz with a minimum sub-block size of 24 MHz consistent with downstream OFDM. The capability to null out subcarriers potentially allows DOCSIS 3.1 to fit in whatever bandwidth is allocated although at a proportionally reduced data capacity. Table 2 details the expected data capacity based on the smallest sub-block size and the range of modulation orders that are most likely to be supported.

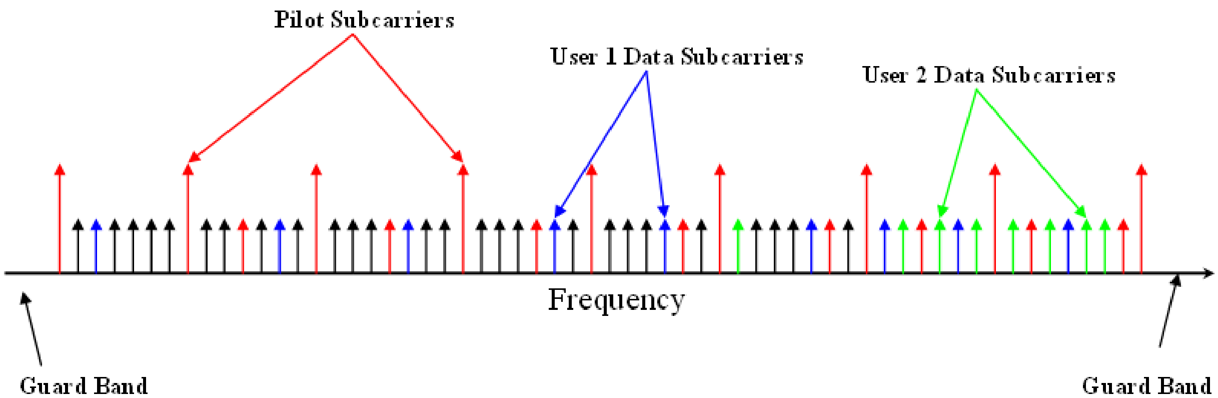

Figure 5 – Example of Multi-user Subcarrier Assignments with OFDMA

OFDM US Data Rate Capacity

QAM	Bits/symbol	FFT Block Sym rate (Msps)	Sub-block size (MHz)	Raw Capacity (Mbps)	Efficiency (estimate)	Delivered Capacity (Mbps)
64	6	96	24	144.00	0.7	100.80
256	8	96	24	192.00	0.7	134.40
1024	10	96	24	240.00	0.7	168.00

Table 2 – Upstream Capacity Calculations for D3.1 OFDM

The upstream DOCSIS 3.1 data rate target of 1 Gbps can be easily achieved using two US FFT blocks of 1024 QAM modulated OFDMA subcarriers. This equates to 192 MHz of upstream RF spectrum. It is anticipated that the working group final upstream allocation for D3.1 will specify a minimum of 200 MHz bandwidth.

Finding room for D3.1 in the current 5 to 42 MHz return path bandwidth is virtually impossible without a mid split expansion. Most cable systems today are using two DOCSIS 3.0 channels (6.4 MHz) and one DOCSIS 2.0 channel (3.2 MHz) to meet existing voice and data rate tier demands. Many operators expect to add a third DOCSIS 3.0 channel within the next year or two. Even with an OFDMA channel minimum sub-block size of 24 MHz there is not enough RF bandwidth available to accommodate D3.1 and maintain the legacy DOCSIS channels. Migrating to an 85 MHz mid split would provide the needed growth room to effectively plan upstream capacity without extensive node segmentation.

INTRODUCING D3.1 INTO LEGACY HFC NETWORKS

When it becomes available, the first applications of DOCSIS 3.1 in legacy 750 MHz to 1 GHz systems will be to raise delivered DS data rate tiers without making any physical changes to the existing HFC plant equipment. The next sections of this paper will review the advantages and limitations of deploying D3.1 in current BW networks. Following sections of the paper will detail the considerations necessary to take full advantage of the data capacity potential of this next generation modulation scheme.

Limitations of the No Touch Approach

HFC networks today are a mix of 750 MHz, 870 MHz, and 1 GHz RF access plants reflecting the system design targets of the individual cable operator. RF bandwidth extensions, analog reclaim, switched digital video, and various digital compression techniques have allowed MSOs to expand the content offerings and IP services they provide while at the same time extending the lifetime of their existing network. Node segmentation provides operators with a minimally disruptive method to significantly increase the delivered bandwidth per subscriber. The cost of the initial primary node segmentation is typically estimated at $20,000 since the majority of the expenses are usually limited to material costs rather than new fiber deployments. [7] By comparison, subsequent node split costs can increase almost exponentially due to fiber construction expenses and when calculated based on the fixed number of subscribers served by a particular node. The appeal of DOCSIS 3.1 in this situation is the potential to increase data capacity per subscriber without doing additional node splitting.

It is expected that data rate growth will continue to be asymmetrical with download speed requirements increasing at a significantly faster rate than upload speeds. Downstream RF bandwidth continues to be under pressure today due to the steady increase of HD programming, the popularity of on-demand streaming, and the rapid growth of IP everything. Many MSOs continue to support a large number of analog video channel offerings due to the large CPE conversion cost of migrating to all digital all at once. Others feel that analog is still a positive differentiator to customers comparing cable with competitive satellite and PON providers. Every MSO expects to migrate to all digital carriage at some point in the future, but the projected timelines vary per operator from within the next 12 months to nearly ten years out. Finding open RF bandwidth is already a challenge in most cable operator networks. Reclaiming channels in order to deploy DOCSIS 3.1 without major disruption to the existing physical plant will take careful planning.

In legacy brownfield networks active element gain, tilt, and spacing along the coaxial access path has been set to optimize the bandwidth and cascade depth reach of the system. For 870 MHz systems and particularly for 750 MHz systems with analog plus QAM channel loads, there are only a limited number of open channels available. In most cases, adding a new OFDM channel block on these networks can only be accomplished by reclaiming RF spectrum from the existing analog or digital portions of the channel map.

In the case of 1 GHz networks, few if any are fully loaded today. DOCSIS 3.1 could take full advantage of this available channel space. In many designs, systems with longer amplifier cascades have a buildup of cable and passive losses along with response flatness issues at the high end of the spectrum preventing acceptable BER / MER performance. OFDM could help these 1 GHz system operators to reclaim this lost bandwidth as illustrated previously in Figure 3.

A new OFDM channel block can be placed anywhere in the downstream spectrum but the most likely location for these subcarriers will be above the existing broadcast and narrowcast channel loading but within the upper band edge of the system. The robustness of OFDM will allow operators to reclaim previously unusable channel space at the upper limit of this RF spectrum. The data capacity increase due to this legacy no touch scenario will be limited only by the amount of bandwidth that is dedicated to DOCSIS 3.1 modulation.

Similarly, the current 5 to 42 MHz return bandwidth that dominates in all North American cable MSO networks is already rapidly approaching its capacity limit. DOCSIS 3.0 and node segmentation have kept cable operators just ahead of the curve but the existing RF bandwidth limits the upstream to just under 100 Mbps using QAM 64 modulation. Low frequency ingress and impulse noise further reduce the usable portion of this narrow allocated spectrum. Many cable systems today load the upstream with two 6.4 MHz DOCSIS 3.0 channels and one 3.2 MHz DOCSIS 2.0 channel. This covers their highest advertised data rate tier plus VoIP phone service but leaves very little spectrum for a new D3.1 sub band. Transitioning to higher order modulation would only provide a short lived incremental increase in capacity.

Extending the Life of the Brownfield HFC Network

The compound annual growth rate of upstream data usage commonly reported at 10 to 12% is consistently lower compared to downstream rates. In spite of this lower growth rate and the benefits of node segmentation the current upstream band cannot meet anticipated forward looking capacity requirements due to the limited 5 to 42 MHz RF return bandwidth. Even with DOCSIS 3.0 channel bonding the peak data rate is restricted to roughly 100 Mbps.

To truly extend the life of the upstream plant and reduce the urgency of node splitting, an increase in the allocated RF bandwidth is needed. As a result, many system operators are now planning 85 MHz mid split trials in 2013. The 85 MHz mid split or commonly referenced "N-split" return bandwidth allocation was defined as part of the DOCSIS 3.0 standard. The shift to 85 MHz would nearly triple the amount of clean spectrum available in the return band with only a small reduction in the forward path bandwidth.

Cable system operators initiating mid split band shifts have in almost every case already moved to all-digital QAM transport. An 85 MHz return band will allow cable operators to maintain existing DOCSIS 2.0 and 3.0 CPE while also providing over 40 MHz of bandwidth for a new D3.1 upstream sub-band.

The mid-split migration will require changes to the diplex filters and return path gain stages located in the node and amplifier E-pac modules. Addressing downstream bandwidth improvements at the same time as the mid split migration would provide a one touch opportunity to extend the data capacity of both downstream and upstream. For 750 and 870 MHz systems this could be as straightforward as changing out the E-pac with a 1 GHz capable version. The major equipment manufacturers have all consolidated their laser, node, and RF amplifier product offerings to 1 GHz designs that drop into existing housings. With the proper padding and equalization these 1 GHz actives can maintain legacy 750 and 870 MHz system performance and provide a future migration path when additional frequency bandwidth is needed.

Increasing the upstream bandwidth beyond 85 MHz could have significant design and cost impacts to legacy brownfield HFC networks. The final determination of a high split frequency plan for DOCSIS 3.1 is still being debated by the CableLabs D3.1 working group. The upstream high split band edge is expected be specified at or near 200 MHz with an appropriately narrow guard band between US and DS that balances the potential impact to CPE cost against a significant reduction in the number of revenue generating DS channels.

GOING BEYOND 1 GHz NETWORKS

The main drivers for increasing the HFC RF plant frequency bandwidth beyond the current 1 GHz DOCSIS 3.0 spec limit are expanding upstream data capacity and at the same time preserving the total existing downstream bandwidth. In order to reach the 10 Gbps target goal of DOCSIS 3.1 or EPoC a total RF bandwidth exceeding the current 946 MHz allocated to downstream channels in a 1 GHz RF plant is required. When combined with the bandwidth needed to maintain legacy broadcast and narrowcast video, phone, and D3.0 data services the total RF spectrum estimates range from 1.5 GHz to 1.7 GHz. A second driver also related to DOCSIS 3.1 is the potential expansion of the upstream to 200 MHz in order to achieve the target goal of 1 Gbps data rates. The only way to accommodate 200 MHz of new return path spectrum is to cannibalize downstream bandwidth. To preserve the current downstream bandwidth most cable operators prefer to shift the downstream upper band edge to 1.2 GHz.

Developers of new HFC plant equipment are already at work planning designs that will support the eventual introduction of DOCSIS 3.1 including extended bandwidth optical and RF plant actives. Whenever outside plant changes are considered the impact on new builds is much different than migrating an existing system where amplifier spacing's, powering, and signal distribution have been pre-determined. The following sections will detail the various changes and challenges of expanding the HFC plant beyond 1 GHz. Since the DOCSIS 3.1 and EPoC working groups have adopted OFDM transport over coax as the central anchor point for their respective efforts, the principles discussed in these next sections are applicable to both technologies.

Extending the Coaxial Network Bandwidth

The first elements that need to be considered for any outside plant migration are the coaxial and RF passive devices. Trunk, access, and drop coax attenuation versus frequency data is readily available from the manufacturers. Figure 6 shows estimated levels and coaxial loss budget information for a simplified N+3 downstream design extended to 1.2 GHz.

Taps and passives with 1 GHz bandwidth have been available for several years and are now ubiquitous across every MSO network. The main line insertion loss and tap port loss of these devices is well behaved across the specified bandwidth making performance estimates easier to generate.

Above 1 GHz the tap port attenuation and thru loss tilt increases substantially with increasing frequency. Cascaded insertion losses from a typical 5 to 7 tap string combined with highly tilted tap port loss further reduce end of line signal levels by an additional 4 to 6 dB at 1.2 GHz. Figures 7 and 8 show examples of the main line insertion loss, return loss, and tap port response for a 1 GHz 14 dB tap plotted from 1 MHz to 1.5 GHz. While the claimed advantages of OFDM modulated channels make it feasible to operate in this imperfect frequency response portion of the spectrum, the high cascade loss budget may overwhelm the available signal level beyond the first few taps in the string.

The example of Figure 6 also illustrates the challenges that many cable operators encounter even with a 1 GHz network deployment. The end of line modem input levels for 1 GHz and higher in this model are at the lower limit of the specified range for most CPE devices. Additional losses due to plant seasonality variations, in-home issues due to customer wiring, etc. will further decrease the received signal levels below the threshold.

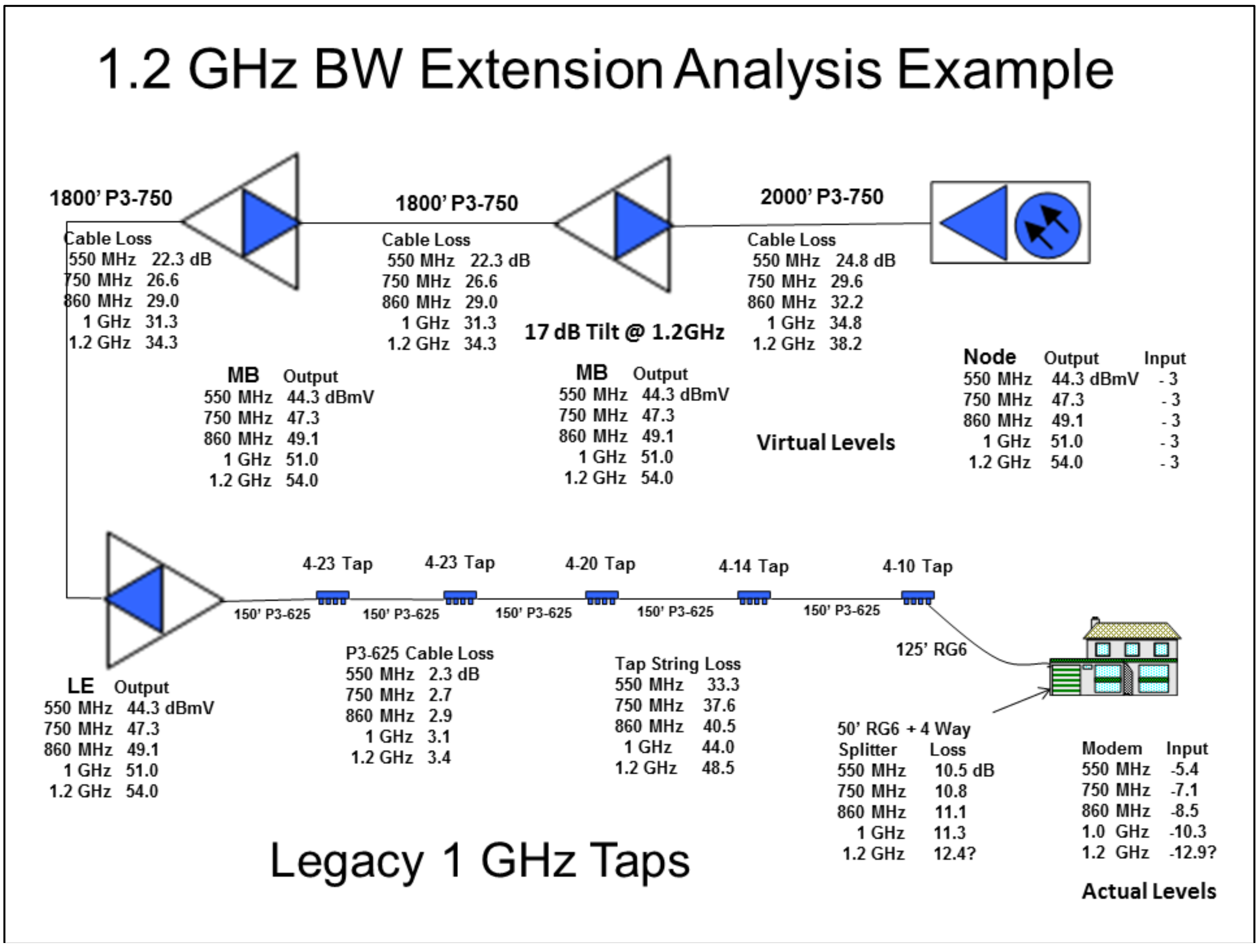

Figure 6 – Design Example of a 1.2 GHz Bandwidth Expansion

Extended bandwidth taps and passives are beginning to appear on the market in response to interest in expanding above 1 GHz. In most cases these devices simply provide a more controlled flatness and return loss response up to 1.2 GHz but insertion losses are not improved. New innovations have demonstrated performance capability up to 1.8 GHz. These devices have also shown improved insertion loss performance compared to legacy 1 GHz passives which could be used to gain back 1 to 2 dB of SNR margin or compensate for other losses in an existing 1 GHz network.

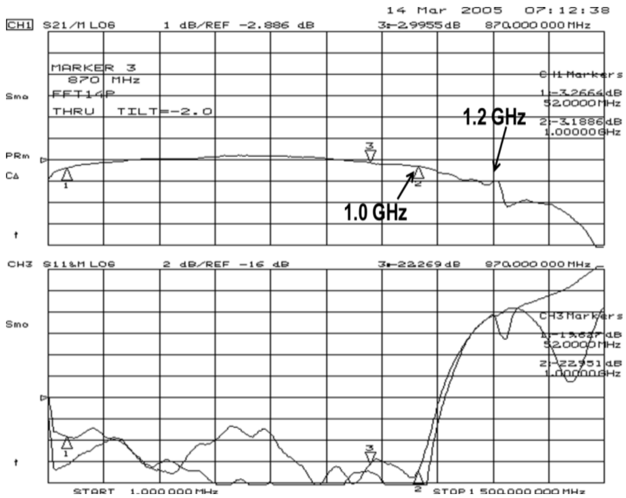

Figure 7 - 14 dB 1GHz Tap Thru Loss and RTN Loss

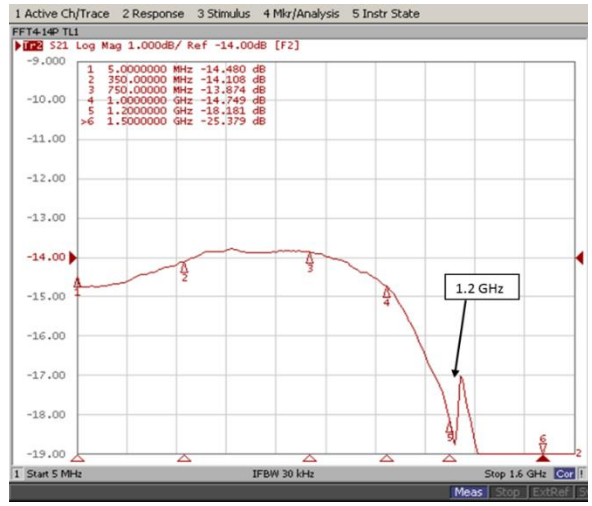

Figure 8 – 14 dB 1GHz Tap Port Loss

Optical Headend Laser Transmitters

Every optical transmitter intended for use in HFC downstream access links today is designed with a DFB laser at its core. These lasers have been linearized either by pre-distortion techniques or through the use of an on chip or external modulator. The RF bandwidth response of the lasers and modulators used in HFC applications are typically 3 GHz or higher. The actual usable bandwidth is determined by the RF amplifier driver stages at the input to the laser or the various laser package parasitics that limit or disturb the broadband frequency response. All laser transmitters capable of analog video loading should also be able to transport OFDM channel blocks.

HFC analog lasers have a fixed optical modulation index (OMI) typically in the range of 22% to 26% for a 1 GHz transmitter. Increasing the loading by 200 MHz decreases the OMI per channel. The effect of this added loading on CNR is a reduction of 0.6 dB per channel. This assumes there is no change in the upstream bandwidth. The reduced C/N could be minimized if the total loading on the laser was only shifted in frequency and kept constant at a 1 GHz level. This scenario would occur if the upstream was expanded to a 200 MHz high split and the 1 GHz downstream loading was merely shifted up to 1.2 GHz.

Creating optical headend laser transmitters that provide a 1.2 GHz flat RF response is a relatively straightforward task of modifying the RF driver stages to assure that they have the bandwidth and linear output capability to drive the laser or modulator. Transmitters designed for this extended bandwidth have already been displayed by a number of vendors. Distortion performance measurements for these extended bandwidth devices is still somewhat variable depending on the equipment used and skill of the vendor.

DOCSIS 3.1 and EPoC modulation format specifications are still in development so signal generation and test equipment availability is still a few years away. Most the specialized QAM generation and test equipment for cable applications on the market today is limited to 1 GHz total bandwidth with the exception of a few high end lab grade instruments. Digital channel loading above 1 GHz can be simulated by the use of noise blocks or by up converting a band of lower frequency QAM carriers. These methods allow basic noise power ratio measurements with the appropriate channel filters. MER and BER testing is possible using up converted QAM channels. The upconverter phase noise is especially critical to making these digital measurements accurately.

Laser transmitters with bandwidths higher than 1.2 GHz are also possible. Many manufacturers have created versions covering various frequency bandwidths up to 3 GHz in order to support satellite and military applications. The drawback for analog lasers as the RF bandwidth becomes wideband is the reduction in carrier to noise which will ultimately restrict the maximum link reach of the system.

Nodes and Actives

Increasing the bandwidth and channel load has the highest impact on the RF active components in the system. Output power level, gain, power consumption, thermal dissipation, path isolation, and numerous other design considerations must be addressed. A primary driver for all network migrations is backward compatibility with the existing deployed network. This places additional pressure on the performance of the expanded BW actives since key factors such as DC power, mechanical housing size, and input / output levels are set by the prior legacy design.

In order to migrate an existing system to a 1.2 GHz network, RF gain and output levels must be increased to maintain legacy 1 GHz performance and overcome the higher cable loss budget at 1.2 GHz. Expanding the channel loading to 1.2 GHz effectively extends the current 14 dB RF amplifier output tilt line common for 1 GHz systems to 17 dB. The added loading increases the output level requirement for the node and each amplifier by 3 dB. The design example diagram in Figure 6 shows the digital channel power (virtual) needed for each station in the cascade. The higher output at 1.2 GHz and raised tilt level increases the power load of the digital channels by approximately 3.5 dB. This will increase the CIN distortion generated by the digital channels which primarily impact the channels in the lower part of the frequency band.

GaAs technology power doubled gain blocks are not capable of supporting the higher output levels needed for 1.2 GHz channel loading. Initial testing using Gallium Nitride (GaN) devices which have been introduced in a number of node and amplifier platforms over the past three years show adequate performance, assuming the digital derate remains at -6 dB referenced to the virtual analog level, but with very little excess margin. The advantage of GaN is its improved output power capability. This is achieved primarily as a result of dramatically reduced thermal resistance compared to GaAs allowing higher output power without increased die temperatures. Another key difference is the higher voltage capability of GaN. GaN amplifier technology was initially developed for high voltage operation applications such as satellite transponders and terrestrial base stations. Cable amplifiers have been designed for 24 volt operation since the first silicon hybrids introduced in the early 1970's. Modified higher voltage power supplies would provide the potential for additional output capability.

The major challenge in the migration path to 1.2 GHz or even higher bandwidths is station gain. Increasing the amplifier tilt is necessary to overcome the increased cable and passive losses associated with higher bandwidth operation. Along with higher output capability, increased gain would normally allow amplifiers to hold their current locations when migrating to the higher frequency bandwidth. The difficulty is that existing cable networks have migrated several times over the past 20 years as a result of bandwidth capacity drivers and new technology innovations. In each case the internal circuitry modules or E-pacs have been updated but the strand mount housing has remained in place. The typical internal gain stages in a 1 GHz amplifier today total up to well over 50 dB. Some of this gain is lost to internal filter attenuation, equalization boards, splitters for multiple outputs, test points, and other necessary functions. The additional gain needed for a 1.2 GHz migration to hold locations in a brownfield design is an estimated 4 to 6 dB depending on the increased losses of new interstage components for above 1 GHz operation.

This is further complicated due to the potential expansion of the upstream bandwidth to 200 MHz for D3.1 and EPoC 1 Gbps data capacity improvements which will require an estimated 4 to 5 dB increase in return path gain. The combination of forward and reverse gain increases will make it extremely difficult to maintain path isolation and stability in the current amplifier housings. This makes it unlikely that traditional 6 deep cascade brownfield networks can be migrated beyond 1 GHz without a major re-design effort.

For new build applications, amplifier spacing's can be set to match the achievable stable gain of 1.2 GHz actives.

<u>Upstream Expansion beyond 42 MHz</u>

With a few component changes most deployed amplifiers and nodes can be migrated to an 85 MHz mid split. These changes include the diplex filters and any high pass cut off filters added to the upstream signal path in order to improve path isolation within the amplifier. Depending on the age of the unit, in some cases the return path hybrid may also need to be replaced. It is not advisable to make these changes in the field so the migration process will include swapping out the amplifier E-pac module with one that has been converted and bench tested to assure proper operation. With appropriate care and controls, lab testing has been done demonstrating that it is possible to hot swap these modules lowering the manpower costs and system downtime during the conversion.

In a high split expansion to 200 MHz or above several additional factors will need to be addressed. First among these is the increased coax loss plus high value tap attenuation that impact upstream modem levels reaching the first active. Figure 9 illustrates the cable losses for different upstream bandwidths and shows the estimated modem level reaching the first active amplifier. The modem output in this example is based on the current D3.0 four bonded channel level. Four D3.0 channels represent a roughly equivalent bandwidth to the minimum D3.1 sub-block of 24 MHz.

Higher cable loss tilt between 5 and 200 MHz shown in Figure 9 requires equalization to avoid large variations in modem power levels across the frequency band. Loss variation over temperature will also increase significantly and may require the addition of return path AGC within the cascade amplifiers.

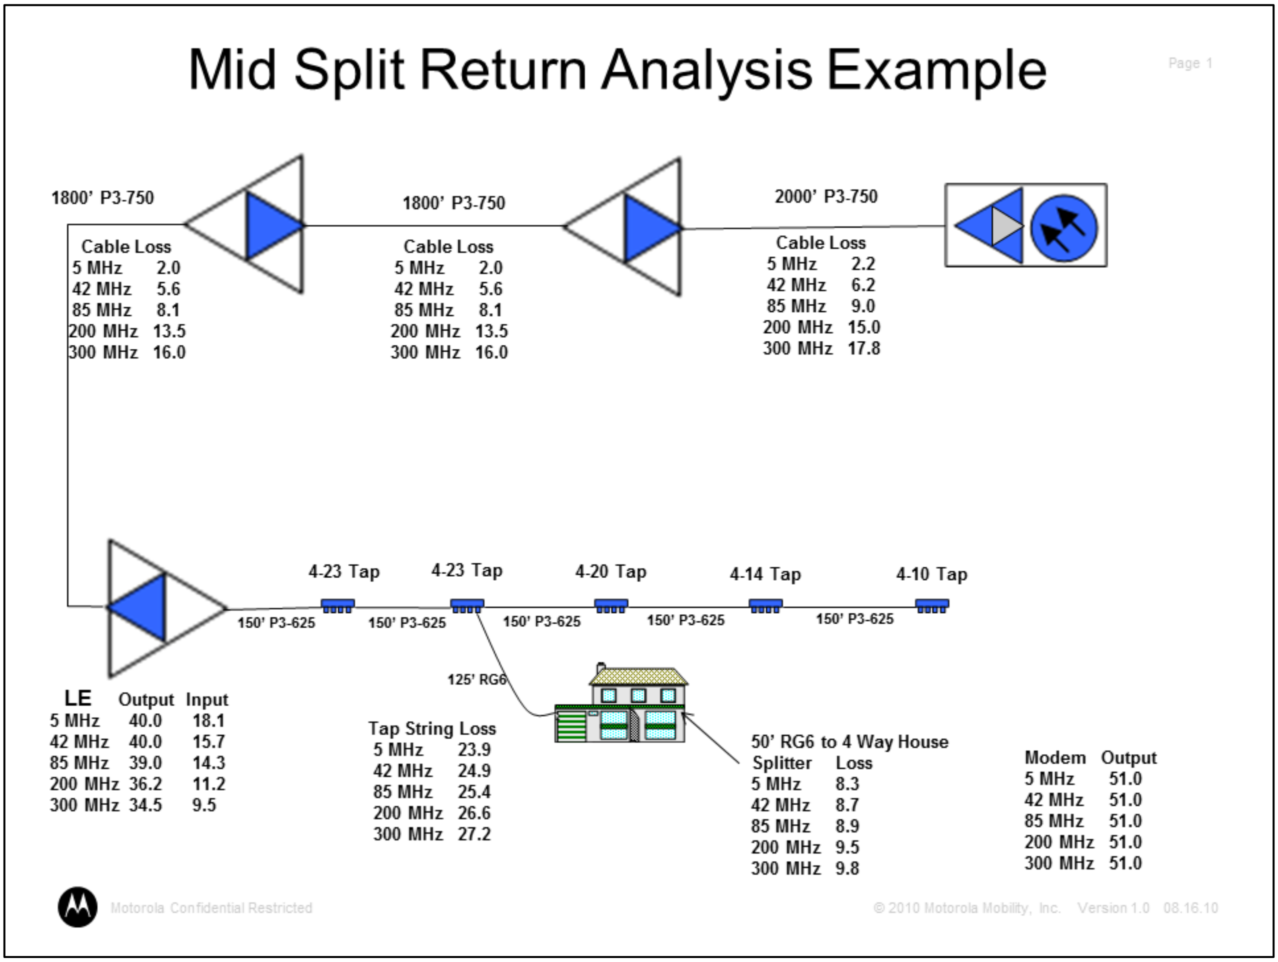

Figure 9 –Design Example for Mid and High Split Upstream Expansions

The SNR delta between D3.0 64-QAM and D3.1 1K-OFDM is 6 to 7 dB. The increase in cable and passives loss as the upstream bandwidth expands to 200 MHz will cause D3.1 channels levels to move closer to the dynamic range noise floor of the upstream laser as shown in Figure 10 unless modem levels are raised higher [(8)]. Since the expected release of the DOCSIS 3.1 specification is still several months away it is unknown what output levels these next generation modems will achieve. This remains a critical issue in the ultimate performance of DOCSIS 3.1.

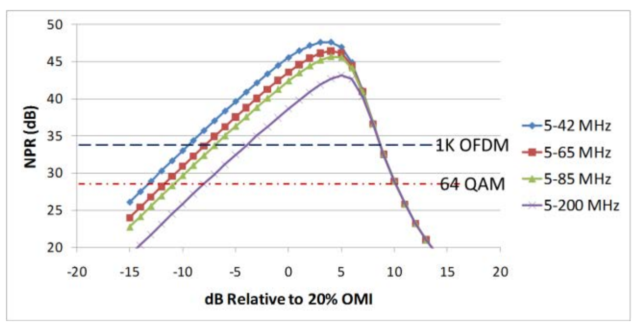

Figure 10 – US NPR for various BW splits

Analog vs. Digital Return

Figure 11 shows measured typical NPR performance of a 2 mw DFB-RPR return of nominal link length. A measured DR system using (post-processed) 10-bits of transport is shown overlaid, in each case using a 65 MHz (European) split.

There is link length dependence for the analog link, and the associated wavelength vs. loss dependence. These variables are not drivers of NPR performance for optical fiber lengths within the digital optical link budget of the DR system, as is commonly the case for HFC applications. Nonetheless, this data confirms the general equivalence of a digital return system achieving a full ten bits of performance to nominally performing higher power DFB returns over average HFC link lengths. It is also apparent how both technologies show comfortable margin to the higher order modulation thresholds shown.

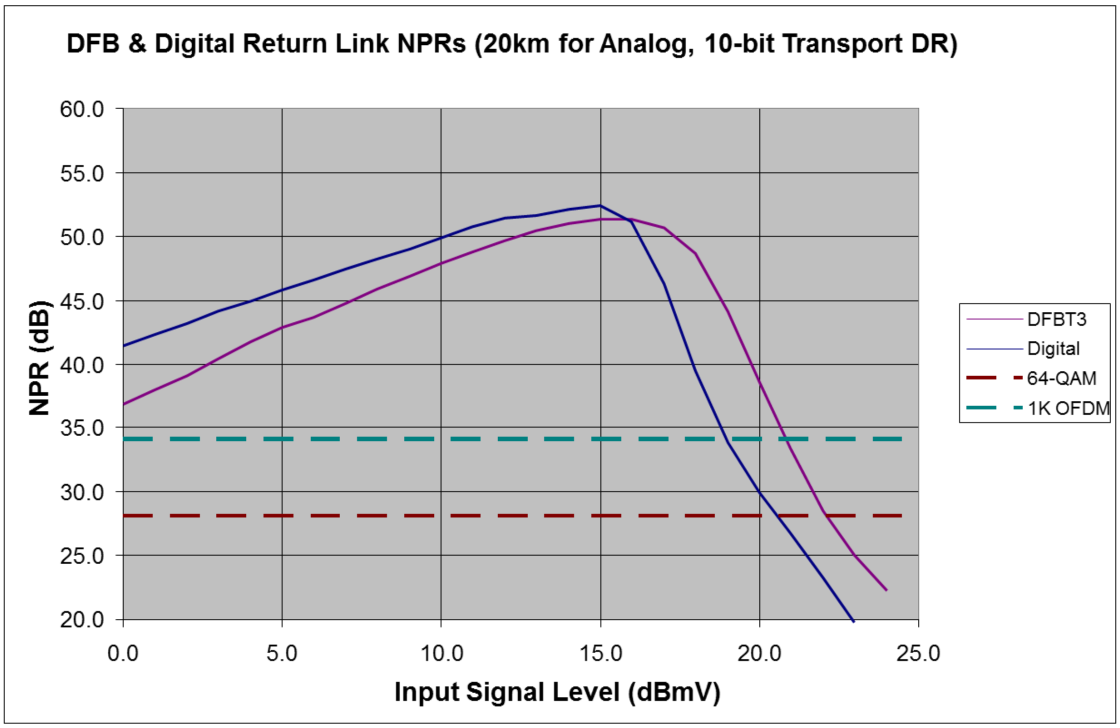

Figure 11 – Typical NPR Performance, Existing DFB and Digital Returns (65 MHz)

Analog return lasers and RF driver stages already accommodate these bandwidth extension options with minimum changes, if any, needed to existing deployed transmitter modules or Hub receivers. For digital return the sampling rate and laser data rate requirements for a typical 2X RF stream transmitter become increasingly difficult and expensive as the bandwidth increases. Table 3 shows the optical line rates resulting from various combinations of A/D resolution and RF upstream bandwidth.

The implication here is that each incremental increase in bandwidth will require a new design iteration replacement of the current DR transmitter / receiver pair. The A/D and laser cost for 200 MHz and higher bandwidth increases dramatically driven by the higher sampling speed and high cost >10 Gbps optics.

Return BW (MHz)	10 bit A/D Sample Rate	12 bit A/D Sample Rate	Laser BW Requirement
5 - 42	1.90 Gbps	2.28 Gbps	2.5 Gbps
5 - 85	3.60 Gbps	4.32 Gbps	4.5 Gbps
5 - 125	5.20 Gbps	6.24 Gbps	8 Gbps
5 - 200	8.40 Gbps	10.08 Gbps	12 Gbps

Table 3 – Digital Return: A/D Resolution, Upstream BW, and Optical Link Bit Rate

CONCLUSIONS

Downstream data rates continue to grow at a 50% compound annual rate and still widely outpace upstream growth. Node splitting is a viable remedy to relieve the increasing data capacity pressures but the cost increases dramatically with each additional layer of segmentation. The introduction of DOCSIS 3.1 in the next two or three years promises to provide cable operators with an alternative tool to incrementally increase data rate capacity to gigabit rates without a major forklift upgrade of the existing HFC plant. The robustness and efficiency of OFDM modulation in conjunction with LDPC coding will enable DOCSIS 3.1 channels to reclaim spectrum that is now impractical to use with prior DOCSIS formats.

The limited spectrum available within the current 5 to 42 MHz return band continues to constrain the peak deliverable data rate and the future ability to effectively use DOCSIS 3.1 to increase upstream capacity when it becomes available. Many cable operators are now planning to trial and potentially deploy 85 MHz mid split systems starting in 2013. To accomplish a mid split migration, the node and amplifier RF modules must be configured with new diplex filters. Migrating to 1 GHz capable downstream modules at the same time as the mid split would achieve a one touch bandwidth capacity expansion that will extend the life of the legacy HFC network for 10 years or more.

HFC frequency extensions beyond 1 GHz are particularly feasible for Greenfield applications. Taps and passive devices allowing future expansion up to 1.7 GHz bandwidth are planned to be available before the end of this year. Optical lasers and GaN amplifier technology has the output capability to support 1.2 GHz networks today and potential optimizations on these devices show promise to further extend bandwidth and reach within a few years. Extending the frequency bandwidth of legacy brownfield networks beyond 1 GHz is going to be much more challenging. The cost of changing every tap faceplate in order to access the higher bandwidth will be the first impediment. Beyond that, the combined gain increases needed for both forward and return signal paths to drive existing 750 MHz spaced housings will require new re-designs to assure path isolation and stability within the node and amplifier. Even then it is not certain that some level of amplifier re-spacing will not be needed.

For the many reasons stated above and detailed in this paper, it appears extremely unlikely that the need for a high split return band and complementary downstream expansion beyond 1.2 GHz will be felt for a considerable number of years. As amplifier cascades significantly shorten and fiber is deployed deeper into the network, the viability of 10 Gbps wideband RF delivery networks will be within the reach of every cable operator.

REFERENCES:

[1] National Broadband Plan available at www.broadband.gov

[2] "FCC Chair Issues Gigabit City Challenge", Jan 18, 2013 (www.lightreading.com)

[3] "On the Evolution of the HFC Network and the DOCSIS CMTS: A Roadmap for the for the 2012 – 2016 Era", Tom Cloonan, SCTE 2008 Cable Tec-Expo Technical Paper

[4] "DOCSIS 3.1 Targets 10-Gig Downstream", October 18, 2012 (www.lightreading.com)

[5] Robert Gallager,"Low-density parity-check codes", IRE Trans. Information Theory, pp. 21-28, Jan 1962.

[6] "Performance evaluation of advanced modulation and channel coding", ReDeSign–217014, 30 Nov 2009

[7] Miguelez, Phil, Fred Slowik, Michael Aviles, "Lowering the Cost of Network Segmentation", SCTE Expo, Orlando, FL, October 17 – 19, 2012

[8] Miguelez, Phil, Dr. Robert Howald, "Digital Due Diligence for the Upstream Toolbox", NCTA Cable Show, Chicago, IL, June 14 – 16, 2011

Acronyms and Abbreviations

AGC	Automatic Gain Control
CCAP	Converged Cable Access Platform
CIN	Composite Intermodulation Noise
CMTS	Cable Modem Termination System
CNU	Coaxial Network Unit
CPE	Consumer Premise Equipment
DOCSIS	Data over Cable Service Interface Specification
DR	Digital Return
DS	Downstream
EPoC	Ethernet Protocol over Coax
FCC	Federal Communications Commission
FDD	Frequency Domain Division
FEC	Forward Error Correction
FFT	Fast Fourier Transform
FTTH	Fiber to the Home
GaAs	Gallium Arsenide
GaN	Gallium Nitride
Gbps	Gigabit per second
GHz	Gigahertz
HFC	Hybrid Fiber Coax
KHz	Kilohertz
LDPC	Low Density Parity Check
Mbps	Megabits per second
MHz	Megahertz
MSO	Multiple Service Operator
NPR	Noise Power Ratio
OFDM	Orthogonal Frequency Division Multiplex
OHE	Optical Headend
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RPR	Return Path Receiver
SCM	Single Carrier Modulation
SNR	Signal to Noise Ratio
TDD	Time Domain Division
US	Upstream

MANAGING THE WI-FI HOME EXPERIENCE

Sandy Howe
ARRIS Inc.

Charles Cheevers
ARRIS Inc.

Abstract

With the average number of WiFi consumer owned connected devices now at 6 devices and growing, cable operators have to deal with a tidal wave of new devices connecting to the home network. These Wi-Fi devices are a mixed blessing – they can help reduce installation costs, but they can drive up customer care calls if proper Quality of Experience is not maintained. Hence, it is becoming increasingly important to ensure that we have the appropriate tools and technologies to manage the evolution to an all WiFi connected home.

First, we will review the trends on Wi-Fi device adoption and the rate at which new devices get added to the home Wi-Fi network. Newly added, heavy hitters for bandwidth or portable devices can present increased difficulties in managing a consistent quality of experience and we will explore how best to optimize their performance.

Next, we will provide customer experience research completed in multiple home environments with different data rates and types. With multiple Wi-Fi devices competing in the home, the consumer experience can be artificially impacted and we will review how this can occur. We will also touch on the effects of supporting close in and far out WiFi devices as they compete for time and resources of the airtime and Access Point to service them.

Once we have defined Wi-Fi trends and measured the quality of experience when not properly managed, we will the review short and long term solutions to address these issues and continue to deliver optimal in-home Wi-Fi experience.

Finally, we will show how, through constant technology updates, we can significantly improve the customer's Quality of Experience. We will analyse the implications and benefits of the following:

- *Hardware Components – 802.11 Technologies, Effects of additional Radios, Spatial Streams, MIMO and Beamforming.*
- *MSO Owned Hardware – Home Gateway and Modems as the manager.*
- *Network and Cloud - Adaptive Bit Rate, Software Monitoring Tools, etc.*

Attendees will leave this presentation with a better understanding of the importance of managing the in home Wi-Fi Network and provide the best Quality of Experience.

ACCESS TO THE PRESENTATION

The audio recording of this topic and the supporting presentation materials will be available at www.thecableshow.com and www.springtechforum.com after the conclusion of the 2013 Spring Technical Forum at The Cable Show 2013.

MORE THAN JUST THE METADATA: HOW CONSUMPTION DATA AND USER PATTERNS HOLD THE KEY TO PERSONALIZING TV

Peter Docherty, Founder and CTO
ThinkAnalytics

Abstract

Cable operators and programmers are looking to content recommendations solutions as a way to offer a more personalized TV experience. First and foremost, next-generation recommendations technologies must be capable of personalizing the primary TV experience. The set-top box experience is critical and often more challenging than performing recommendations across the second screen. This is because delivering second screen content is similar to VOD architecture and traditionally easier to implement in comparison to the primary TV experience, which involves dynamic and linear TV (including EPGs).

These next-generation recommendations systems rely on a combination of robust metadata and a variety of recommendations algorithms, including implicit and explicit data points, direct user feedback such as likes and dislikes, and consumer viewing behavior. But also crucial to these recommendations and navigation technologies is understanding the consumption data and user patterns of linear real-time programming so that operators can maximize their content libraries and add value to existing services. Success hinges on being able to extract valuable and useful learnings from all the noise in the zapping data, and relying on a methodology and approach that is capable of scaling across the demands of real-world operator deployments. For example, a pay-TV provider with 10 million subscribers offering 80,000 programs would need interacting algorithmic technology capable of supporting 800 billion recommendations combinations.

Incorporating these consumer-centric metrics creates an opportunity to understand a consumer's unique preferences, incorporating data such as explicit preferences, implicit preferences, moods, likes, dislikes, etc. Additionally, this kind of data can extend the use of important consumer learnings to any part of an operator's organization, such as marketing, retentions, package configuration, acquisition, and channel marketing.

The paper will delve into the technical specifications and advantages of implementing personalized recommendations techniques that personalize cable TV services on the primary screen and extend that experience across the broader device ecosystem. It will compare and contrast competing approaches, and outline how incorporating consumption patterns and viewer usage patterns improve the overall accuracy of content recommendations. It will look at real-world implementations of how operators today are using these new data sets to truly understand and enrich the subscriber experience.

OVERVIEW

According to the Nielsen Cross-Platform Report (http://www.nielsen.com/us/en/newswire/2013/the-cross-platform-report-how-viewers-watch-time-shifted-programming.html)

Americans spent more than 34 hours per week in front of a TV set in the third quarter of 2012. In addition, Americans spent close to five hours a week on a computer screen, using the Internet and watching video content.

Around the world, the average consumer has 400+ programming channels to choose from yet only watches an average of eight channels.

The sheer volume of content is the root cause of one of the most prevalent misperceptions among TV subscribers - that there's "nothing on to watch." In reality, there are so many choices available that it's virtually impossible for consumers to discover what's really there. The challenge is proactively presenting subscribers with a personalized experience which includes content recommendations that truly reflect the viewing preferences of the individual subscriber, instead of forcing them to work hard to find something they like to watch.

Personalized content recommendations help operators provide higher consumer engagement, promote content and boost revenues. Once deployed, operators can further leverage the content usage patterns and viewership data provided by recommendations engines to help aid in customer acquisition, retention and loyalty.

SEARCH, DISCOVERY AND RECOMMENDATIONS DEFINED

Discovery is the new buzz word, but in reality, it isn't new at all. Relying on discovery alone means that subscribers have to work hard to find new content.

Search and basic discovery solutions assume that subscribers know exactly what they want to watch at any given moment. These approaches assume that viewers know the actor, name of the movie, or the type of show. For example, mood-based content recommendations are just another form of search, and require customers to do the leg work to find which content reflects their "mood."

Search is only useful if the consumer already knows exactly what they want to watch, which is only about 4% (http://advertising.yahoo.com/blogs/events-blog/tv-doesn-t-just-live-tvs-anymore-181752027.html) of the time.

A recent study (http://www.rovicorp.com/insights/index.htm) conducted by Rovi revealed that even when a viewer knows what they want to watch, the vast majority can't find it. Specifically, 72% of those surveyed indicated that they could not find the desired content when searching an on-demand library.

The most successful approach to personalized recommendations is a proactive tool that suggests new content, uniquely created for each viewer, based on their viewing history, likes, dislikes, and moods. These advanced systems expand customers' tastes, pulling from the entire content library, whether that is VOD, live linear TV or other types of content. Personalized recommendations also allow for operators to individualize recommendations to subscribers within a household, not just delivering household level recommendations.

Earlier this year, Cox Communications (http://cox.mediaroom.com/index.php?s=43&item=643) became the first cable operator in the U.S. to launch a new class of personalized TV recommendations as part of its *Trio Program guide*(SM). The guide presents recommendations across live TV and VOD content choices that can be personalized to individualized members of a household, and it's a bellwether of what's to come for TV viewing across the U.S.

Personalized recommendations can be a tool for operators in helping subscribers easily find and discover new content and subsequently, gain more value from their subscription packages. In addition, by adding search features to recommendations the viewer gets the best of both worlds with fully personalized and appropriate recommendations plus the ability to search for specific items of interest and content.

CONTENT RECOMMENDATION TYPES

Basic Vs. Advanced Recommendations

There are multiple types of content recommendations ranging from the most basic applications to the highly advanced. Each type of recommendations has its own use case. Broadly speaking, content recommendations can be grouped into the following categories: Collaborative Filtering, Unique Preference Profile, Mood-Based and Content-to-Content Recommendations. The specific types are outlined in Figure 1 below:

Figure 1: Content Recommendation Types

Social recommendations or collaborative filtering

One of the most common forms of content recommendations is Collaborative Filtering. Collaborative Filtering is a basic type of recommendations that uses a web-based approach and relies on driving the most popular content offerings based on viewing patterns. This is similar to the social recommendations approach, "people who watched this also watched that." The disadvantage, however, is that this technique simply drives the most popular content i.e. "the Harry Potter effect", in that everybody gets offered only the most popular content.

It's more advantageous for operators to ensure that a wider range of content is presented to the user that broadens viewing across both VOD and linear TV. This ensures that consumers get overall value from their subscriptions as well as leading them into other areas of content, including Pay-Per-View and transactional VOD.

Mood-based recommendations

Mood-based recommendations also fall into the basic recommendations category. This form of content recommendations relies on a consumer to identify what mood they are in and select from predefined categories, e.g. "I am in the mood for a dark comedy." Another shortcoming is the categorization of the content itself. Placing content into static categories does not allow for a dynamic content recommendations experience in which new content is being offered on an ongoing basis. Mood-based recommendations are really just another form of search, where the user is doing the up-front work of telling the system what to find for them.

Content-to-content recommendations

Content-to-content recommendations are a basic form of recommendations that relies on matching related content based on program content similarity. This means recommending content that has similar actors, genre, mood, etc.

Unique personalized recommendations

Advanced recommendations use automated techniques such as real-time data-mining, natural language analysis, consumption, ratings and other semantic and analytic techniques to constantly learn the preferred viewing habits of a customer. These forms of recommendations take into account many different factors about the content and the user, adapting instantly to customer feedback and to the constantly evolving content.

Intelligent search and navigation

Intelligent Search means that different results can be offered for different subscribers that happen to search for the same term. These results are based on an understanding of individual subscribers and their unique behavior patterns, likes and dislikes and generate search results that are unique to them.

Intelligent search and navigation features include the ability to use user's behavior and preferences to target searches more effectively.

Most deployments of recommendations today need to move beyond basic recommendations to encapsulate a combination of basic and advanced techniques in order to be successful.

Operators need a broad-based comprehensive search and recommendations solution to drive optimum customer value and satisfaction. Personalization, in addition to social recommendations is fundamental.

Recommendations technology that encompasses a comprehensive range of fully automated analytical techniques will offer operators a long-term recommendations solution with the ability to cope with the variety of scenarios that will be needed in the long-term; thereby avoiding the limitations of a one-size fits all, black-box solution. A multi-faceted approach incorporating basic and advanced recommendations ensures that recommendations will support variety of use cases and be executed across the entire device ecosystem.

METADATA

Content metadata and other associated metadata is a key component in providing accurate and successful content recommendations. Metadata is used to identify key attributes generating an enhanced content description such as – star rating, era, mood, theme, atmospheric, age certificate, classification, etc.

Metadata and content profiling

Advanced recommendations engines should provide a content classification process that spans across multiple metadata sources including TV programming, VOD and third-party content. This classification process should integrate with the operator's existing and available metadata sources such as broadcast TV, time shift TV, VOD, third-party product data, marketing campaign systems, etc.

SUBSCRIBER PROFILING

Implicit/explicit preferences

Both Explicit (manual) and Implicit (learned) preferences are a key part of building subscriber profiles so that the learned preferences can be used to influence the profiles in certain directions.

Explicit preferences include information that the consumer deliberately tells the system through answering a questionnaire or rating content e.g. I like/don't like the following content. Implicit preferences, on the other hand, involve information that the recommendations engine learns from the consumer's behaviors and actions (e.g., watching a program, recording, or downloading, etc. Learned preferences (implicit) are used to capture customer preferences without the customer having to go through a series of forms or questions. A key element of learned preferences is that they enable the capture of more subtle information that will allow the recommendations engine to widen the customer's view and get them out of their narrow view of the programming that they typically watch and the channels they would browse.

Incorporating explicit and implicit preferences enables the recommendations system to become increasingly accurate as it continues to learn and adjust as the consumer's tastes change and evolve.

Customer profiling

Recommendations engines provide support for a number of user profile levels such as anonymous users, individuals, households and communities.

By identifying the individual, greater accuracy in recommendations is achieved as the engine identifies distinct behavioral patterns leading to successful recommendations and targeted marketing. Stored historical recommendations become invaluable pieces of consumer intelligence.

CONTENT AND VIEWING HABITS

It is essential when analyzing consumer and viewer behavior that no matter where subscribers view content that the recommendations experience is consistent and aligned to the operator's objectives.

The growth of viewing platforms provides both the consumer and operator to engage in viewing experiences beyond the Set Top Box and this is increasing by a magnitude of around 3-4 from 2012 to 2016. Sales of connected devices such as tablets and smart phones will grow from approximately 750 million today to over 2.5 billion by 2016. The consumption of content will increase in line. Figure 2 below outlines the types of content and platforms that should be supported in order to deliver advanced and personalized recommendations.

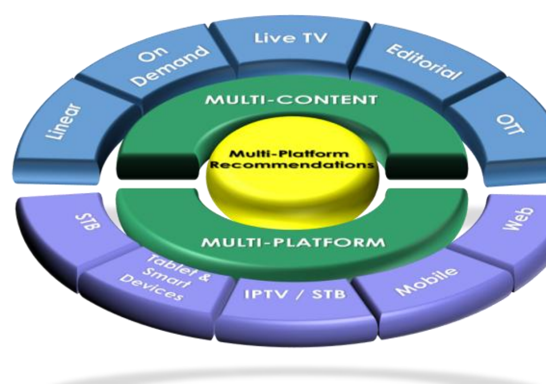

Figure 2: Multi-Platform Recommendations: Content and Device Platforms

Data ownership and storage

The data viewership and usage data is hugely valuable as it reflects not just genre preferences but individual program likes and dislikes. This ongoing data and derived intelligence should be owned and controlled by the operator and can be stored wherever the operator chooses. If not, this highly valuable data about operator's own customers could be owned by a third-party, such as a recommendations engine provider and the value to the operator is significantly diminished.

Broadening consumer tastes based on viewing patterns

To improve customer retention and grow ARPU, it is essential to expose consumers to a wider range of content than they would normally discover for themselves. Where the consumer explicitly declares a wide range of tastes this is relatively simple to achieve, but in most cases (especially during the critical early phase of the relationship with the recommendation system) the customer is cautious and will declare a narrow range of tastes or even none. In response, instead of recommending a default Top 10 list based on collaborative filtering, the recommendations engine can generate a fresh and diverse list of programs every time, using analytic techniques and not on random selection or wholly marketing based manual selections.

In this way, different programs are selected from the live content, including Long Tail surprises or one-offs to intrigue the consumer, and as a result he is encouraged to divulge more and progress his relationship through the crucial start-up phase.

Viewership data as an operator marketing tool

The viewership and content usage data that is gathered for content recommendations can also be used as a real time marketing tool. This can be highly effective with inbound and online campaigns. It allows the marketing message to be precisely targeted based on the very latest data, constantly adapting throughout the interaction with the consumer.

Support for live and on demand content

The preferred content recommendations engines should be able to provide recommendations on-demand and in real-time.

It is much more difficult to develop and manage an intelligent navigation system and recommendations engine capable of handling dynamic, linear EPGs. The sheer scale of the content involved – typically 500 or so channels offering up to 120,000 programs over a two-week period - makes this much more of a challenge, especially for solutions based on collaborative filtering or other single-algorithm/black-box solutions. Live linear or EPG content is constantly updated, meaning that the recommendations engine and intelligent navigation system needs to work much harder than in a VOD environment.

As the subscriber uses the EPG the recommendations engine can be invoked to provide recommended content for that consumer at an individual and/or household level. When the EPGs are updated advanced recommendations engine should process new items, producing and appending new and enhanced content feature classifications for each item.

IMPLEMENTATION AND RESULTS

Operators around the world are operating in mature, highly competitive markets. They want to keep their best customers and drive revenue growth through increasing usage of existing services and creating personalized offers that will lead to increased up-take of these services while at the same time reducing churn in the core customer base.

In one commercial deployment for example, an operator noticed valuable customers leaving to join competitors. These customers did not view offering as a good value and were not able to find ‘good’ content. This caused high attrition rate, poor communication and not engaging the customer at the point of contact.

Using a personalized recommendations system provided insight on how a uniquely intuitive, personalized and intelligent initiative would lead to an increase in their subscriber base while improving their customers’ personal experience with content through the various communication and interactive channels. In order to support this initiative, the operator deployed a recommendations system with ThinkAnalytics to achieve the following objectives:

- Engage and retain existing subscribers
- Entice subscribers to discover more programs
- Show subscribers that “there is something on”
- Raise awareness of the breadth of channels and programs available across the live EPG and VOD platforms
- Make subscribers aware of the range of content and guide them to select more appropriate packages through personalization
- Entice subscribers to take up other value added services and products

As a whole, the customers of ThinkAnalytics as the largest worldwide provider of personalized recommendations deployments have seen significant business benefits in the viewing of recommended content, increased consumption of transactional/paid content and improved customer satisfaction and loyalty, with the recommendations themselves having a positive response from the consumers.

SUMMARY

Research and results from large scale deployments have proven that delivering better TV recommendations can have a profound impact on business results. Fundamentally, any recommendations deployment needs to support the operator’s current objectives and have the inherent flexibility and scalability necessary for the system to evolve. The technology now exists to let operators configure the parameters that best meet their objectives based on analysis of viewing and usage patterns of consumers in order to provide them with a personalized TV experience across platforms. Personalized recommendations can provide a powerful tool to reduce churn, market new or underexposed content, increase on-demand sales, and improve consumer satisfaction.

Our Wi-Fi Future: Technology and Spectrum Enablers

Rob Alderfer, Vikas Sarawat, Kyung Mun
CableLabs

Abstract

Several technology and wireless spectrum developments are on the horizon to help MSOs, and all ISPs, meet the unique needs of wireless video and other growing network demands. This paper reviews these developments and the effect they are likely to have. In the technology space, advances in standards and network optimization techniques will advance network performance and make Wi-Fi more useful for consumers. In addition, new wireless spectrum availability will help to facilitate these advances, and will enable additional capacity and frequency diversity for expanding Wi-Fi use cases.

The goal of this paper is to provide an overview of emerging technologies and new wireless frequencies that may aid Wi-Fi video delivery, how these advances will improve the end-user experience, and what is needed to realize this future vision. We explore how a typical consumer may tangibly benefit from these advances through use cases representing a 'day in the life' progression of Wi-Fi video activity. We do not endeavor a complete technical guide to the implementation of relevant advances in light of their broad scope; rather, our intent is to provide a window into what video over Wi-Fi may look like in the near future.

INTRODUCTION

Video data traffic comprises 65% of downstream network traffic at peak times, and the share is growing[i] as total traffic volume continues to expand,[ii] making it among the most significant factors in ISP network planning. (See Figure 1.) In addition, consumers are connecting more wireless devices than ever before to consume video. Tablet sales have grown by over seven times in the last three years,[iii] are forecasted to grow another three times by 2016.[iv] 90% of tablets only use Wi-Fi, and have no mobile connectivity.[v]

While broadband consumers rely on the cost and performance advantages of the fixed network, particularly for high-bandwidth applications like video, they also value the ability to be 'untethered'. These trends pose challenges for multiple-system operators (MSOs) [vi] as they work to meet growing consumer demand for wireless video over Wi-Fi networks.

Cable operators are continuing to increase network capacity and deploy Wi-Fi access points inside and outside of homes to meet growing consumer demand. The CableWi-Fi initiative, announced last year, enables Wi-Fi access for Cable customers across over 120,000 access points [vii] (and growing) deployed by the five largest MSOs.

This effort is a significant development in Wi-Fi networking, which has historically been ad-hoc and driven by fixed broadband consumers without broader coordination. Access point deployments outside of the home in 'main street' settings, small businesses, sports venues, and other locations represent a new approach to Wi-Fi access – one that is coordinated and provides reliable capacity in most areas that consumers spend their time. The success of this effort is driven by utilizing fixed network assets to extend the value of a product with high consumer demand – high speed data.[viii]

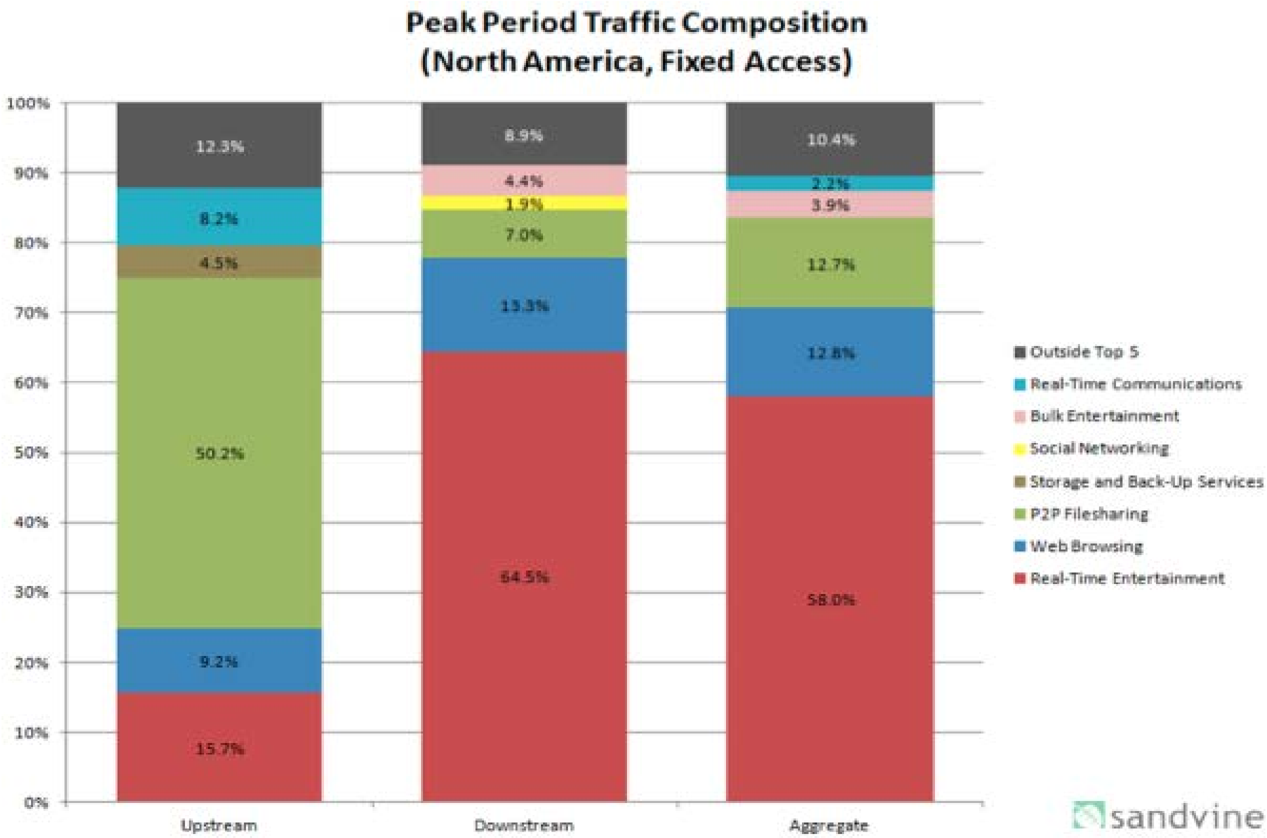

Figure 1: Sandvine Estimates of North American Peak-Period Fixed Traffic Characteristics[ix]

Coincident with this evolution in Wi-Fi networking are consumer expectations of ubiquitous wireless access. Customers expect their connected wireless devices to 'just work'. This desire for seamless connectivity and simplicity is not always today's reality in Wi-Fi world, however. Users must manually find, select, and authenticate to trusted networks, in most cases. Moreover, access point range is limited, exacerbating this complexity for Wi-Fi users on the move. Congestion is expected in crowded areas. These issues become most evident in video applications, which require significant bandwidth, and connection quality impairments are often noticeable in demanding environments.

MSO-delivered Wi-Fi is beginning to bring customer's experiences in line with their preferences. The use of the common "CableWi-Fi" SSID simplifies network selection and authentication, and dense deployment of access points help to limit congestion. While helpful, these steps are work-arounds within the current parameters of Wi-Fi.

On the horizon are several developments in Wi-Fi technology and spectrum access that can fundamentally re-shape the user experience, particularly for video services, which generally require significant and consistent throughput across diverse use environments.

New core standards and spectrum will enable throughput in excess of 1 Gbps, an improvement of nearly 3X in performance over today's Wi-Fi networks. The next-gen television experience will enable wireless content delivery, simplifying in-home networking. Outside of the home, new Wi-Fi technology will enable automatic user authentication, enhancing the utility of Wi-Fi to consumers on the go. Advances in network management will help to ensure consistent user experiences and minimize congestion. And availability of new wireless spectrum will complement these advances, enhancing capacity, reliability, and coverage.

Together, this progress will meaningfully improve the consumer's Wi-Fi experience, helping to ensure that it remains the wireless network of choice. To demonstrate what these improvements will mean in concrete terms, we will track a typical Wi-Fi use case, following a hypothetical consumer through a morning of Wi-Fi use. From when our consumer wakes up in the morning to the time he arrives at work, new developments in Wi-Fi will clearly improve his (or her) experience.

This next phase of Wi-Fi is not a given, however. Several regulatory and technology changes are needed to realize this vision. In particular, regulatory dependencies are uncertain, and will require sustained engagement on the part of the Wi-Fi community. Our hypothetical consumer – and real consumers everywhere - demand it.

TECHNOLOGY DEVELOPMENTS

Several advances in Wi-Fi technology are on the horizon and have the potential to significantly advance Wi-Fi user experiences. These advances, driven by improvements to core standards, network discovery and selection, wireless delivery to the TV, authentication protocols, and network management, including radio resource management, together can help to re-shape Wi-Fi from a manual, inconsistent user experience to one of simplicity and seamless connectivity, especially for video applications.

Core Standards

IEEE 802.11ac is a fifth-generation Wi-Fi networking technology that promises about three times the throughput capacity than the previous generation 802.11n. (See Figure 2 for the evolution of Wi-Fi standards.)

802.11ac contains several characteristics that enable throughput gains over current Wi-Fi technology. One feature is spectrum use: Whereas 802.11n uses both the 2.4 GHz and 5 GHz frequency bands, 802.11ac uses only the 5 GHz band. 5 GHz is a 'cleaner' environment than 2.4 GHz, which is populated with the majority of Wi-Fi devices in use today, along with other wireless devices and services that may cause interference. By using the 5 GHz band, 802.11ac will enable less interference to and contention among Wi-Fi users, enabling higher potential performance. However, regulatory complexity in 5 GHz terms of access may complicate implementation of 802.11ac. These issues will be explored in a later section.

Of course, a cleaner operating environment does not alone guarantee higher performance. 802.11ac also uses wider channels of 80 MHz and 160 MHz, up from the 20 MHz and 40 MHz channels of 802.11n. Expanding channel bandwidth by 4 times the current standards yields the potential to expand theoretical performance by fourfold, holding other factors constant.

In addition, 802.11ac enables greater spectral efficiency and reliability. This is achieved through a doubling of potential multiple-input, multiple-output (MIMO) spatial streams, from 4x4 in 802.11n to 8x8, as well as a more sophisticated modulation and coding scheme, up from 64 QAM to 256 QAM.

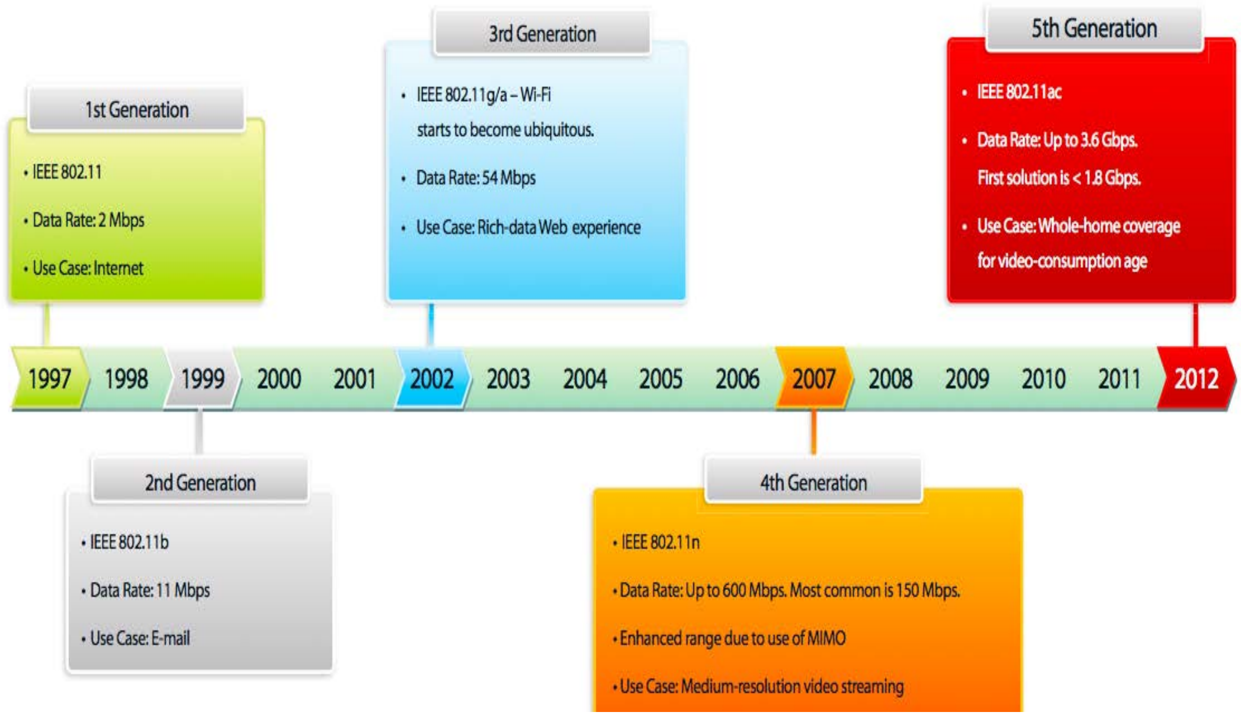

Figure 2: Wi-Fi Core Standards Evolution (Broadcom)[x]

The benefits of these improvements will be aided by standard beamforming in 802.11ac, enabling directional signaling and greater interference avoidance. Beamforming is not widely deployed in 802.11n equipment; its standard adoption will enable enhanced throughput and coverage range.

While 802.11n with MIMO supports enough raw bandwidth for wireless video streaming, these additional features in 802.11ac should make it better-suited to high-quality wireless video delivery.

Wireless to the TV

802.11ac and other technology developments will help enable wireless content delivery to the television. Currently, if a customer wishes to receive video content wirelessly, a common approach is to attach a Wi-Fi-capable third-party box, such as a Roku streaming player, to view IP-delivered content. If the TV has a Wi-Fi receiver, it may connect directly. However, for most consumers today, viewing linear, on-demand, or DVR video content requires a wired connection to the TV from the relevant Cable box, which also requires a connection to a coax output. Several technology advances are likely to improve wireless delivery of video content to the TV.

802.11ad, another IEEE wireless communication standard, enables a short-range direct link of up to 7 Gbps using the 60 GHz band. This standard, also known as "WiGig",[xi] is unlikely to serve as a general-purpose Wi-Fi technology due to path and propagation losses of its high-band spectrum, but may prove to be a replacement for a HDMI cable, enabling wireless delivery of a greater range of video content to the TV.

Equipment using this standard is currently under development, and the Wi-Fi Alliance is expecting to launch certification later this year.

While a robust network is a must for reliable video delivery, users also expect the system to be easy to use, with automated connection and data and video transport. DLNA, Airplay and Miracast are a few technologies that help with this.

Airplay is Apple-proprietary and allows various Apple devices on a LAN to find, connect and stream to each other. However, it is not interoperable with non-Apple products. For example, Airplay enables users to wirelessly stream video from their iPhone to an Apple TV, but an Android device could not connect to Apple TV through Airplay.

Unlike Airplay, Digital Living Network Alliance (DLNA) standards define interoperability guidelines for sharing of digital media across consumer devices of any vendor. DLNA uses Universal Plug and Play (UPnP) for media management, discovery and control. There are thousands of devices from a multitude of vendors that are "DLNA Certified". Some estimates indicate that there are hundreds of millions of DLNA enabled devices have been installed in customer homes,[xii] and many Cable operators use it as a basis for in-home video streaming. Additionally, DLNA is transport-independent and can be used with either wireless or wired network.

In addition, peer-to-peer video delivery will be enabled through a technology known as Miracast. Miracast is a Wi-Fi Alliance certification program for the Wi-Fi Display specification, and like DLNA and Airplay it defines device and service discovery, including the full application protocol stack for the delivery of multimedia. Miracast, unlike DLNA, is transport-dependent and only works with Wi-Fi Tunnel Direct Link Setup (TDLS) and Wi-Fi Direct. Standardized device-to-device video delivery will enable innovative use cases, such as viewing tablet content on the TV, without the need for a physical connection and while minimizing congestion associated with broader networking. Wide adoption of this technology, enabled through the Wi-Fi Alliance certification program, will enable content delivery across device types.

Authentication & Roaming

Currently, connecting to a new access point requires a user to manually search networks, select a trusted SSID, and authenticate. This process complicates Wi-Fi use for consumers and deters the use of Wi-Fi relative to more seamless network access and roaming on mobile networks.

The incorporation of new technologies under a Wi-Fi Alliance certification program known as Passpoint will simplify Wi-Fi network connection. Key elements of the Passpoint program include the Hotspot 2.0 and 802.11u standards, which enable network discovery and selection by communicating identifying information between the access point and client device prior to authentication. Prior to connection, supported client devices can obtain information from the access point that is necessary for network selection, such as operator policy, authentication protocol methods, and other relevant information.

In addition to automatic network discovery and selection, Passpoint also supports several authentication protocols that may ease secure network connection for users. In addition to a username/password process, which is similar to SSID authentication today, Passpoint also supports SIM-based (e.g. EAP-SIM) authentication so that operators can use the same credentials for both mobile and Wi-

Fi. In addition, a certificate-based authentication is possible using technology such as EAP-TLS, which allows operators to authenticate using device certificates.

In addition, Passpoint standardizes WPA2-Enterprise to enhance the level of security in Wi-Fi networks. Later releases of Passpoint will enable operator policy provisioning to manage business relationships, as well as online signup of new subscribers.

802.11r is also important, as it helps enable fast roaming of clients between closely located Wi-Fi Access Points. 802.11r achieves fast roaming by combining key negotiation and requests for wireless resources, targeting handoff time of less than 50 milliseconds.

Together, these advances will make the manual, SSID-based Wi-Fi connection process a thing of the past. Consumers will find it much easier to connect to new Wi-Fi Access Points through what will seem to them like an automatic process. This will enhance the utility of Wi-Fi, particularly for users that are on the move, and makes video applications more feasible when moving between Wi-Fi Access Points.

Quality of Service

Today, Wi-Fi quality of service can vary widely depending on the use environment. Within the home, service quality has historically been enabled through low contention ratios enabled by access point density relative to wireless devices, but as the number of devices in the home grows, along with wireless video traffic, this ad-hoc approach may become less effective. Outside of the home, service quality quickly suffers as access points become crowded. However, new techniques in radio resource management (RRM) and network optimization will enable a more intelligent approach to Wi-Fi service quality.

Many quality of service issues originate in today's simple Wi-Fi connection protocol – the access point with the strongest signal will serve as the default connection, and client devices may stay connected to the same access point even as the signal diminishes and others would provide a better experience. Advances in radio resource management through technologies like 802.11k will enable client devices to select the best available network based on a broader range of environmental information. For example, 802.11k may enable a client device to select an access point with a weaker signal if the access point with the strongest signal is handling a heavier amount of traffic – thus enabling higher throughput to the client device. On a network basis, this approach will ensure appropriate distribution of traffic across access points and spectrum channels, providing a more consistent experience overall.

In addition, broader Wi-Fi network optimization techniques (SON) will aid service quality. For example, throughput may be improved by dynamically managing access point transmit power when two geographically-proximate APs are operating on the same channel. Or, power may be varied depending on relative traffic loads. Latency-dependent applications like video can be provided higher priority on the network through Wi-Fi Multimedia technology. [xiii] Also, beamforming and MIMO approaches can be implemented to manage access point traffic loading. These approaches will also aid in handoff between access points and determining optimal frequency band utilization.

Together, these solutions represent a migration toward network-based Wi-Fi service quality, which will enable more consistent experiences across devices on the network. In addition, software used to implement these optimization techniques can also transmit network performance information to the operator, including maintenance needs, with the effect of minimizing costs and increasing network uptime.

Technology Roadmap

We have described a number of advances in Wi-Fi technology that will improve the user experience in wireless video and other applications. However, realizing these improvements is dependent on broad adoption. Broadly speaking, the technology described here must be adopted by the consumer electronics community, both at the access point and in consumer devices.

Controller-based access point deployments, incorporating optimization technology described here, will provide the network-level control and data needed to move from an ad-hoc Wi-Fi world to one that enables a more consistent user experience.

Equally important is adoption of enabling technology in consumer devices, such as Passpoint, that will make the use of Wi-Fi a more seamless experience for consumers.

Mobile devices are often provisioned with specifications determined by mobile carriers, particularly in markets like the United States where carriers subsidize most mobile devices. In unsubsidized devices, including tablets and many other product categories that do not rely primarily on the mobile network, consumer electronics manufacturers should find it in their interest to optimize the Wi-Fi experience for their customers. In light of the heavy reliance of mobile devices on offloading to the fixed network via Wi-Fi, ensuring the best user experience over Wi-Fi should serve the entire wireless broadband ecosystem, and adoption of enabling technology should be a priority across all Wi-Fi capable devices.

SPECTRUM DEVELOPMENTS

Wireless spectrum is also important to improving the Wi-Fi experience. The 2.4 GHz band remains the most widely-used unlicensed band today; however, heavy usage by WLAN devices as well as baby monitors, cordless phones, microwaves, and other devices has given rise to congestion. The 5 GHz band is also available for Wi-Fi and has plentiful bandwidth, but divergent terms of access complicate its use. For example, 80% of the channels are required to transmit at a lower power than the 2.4 GHz band. Further, the 5 GHz band experiences greater propagation and path loss than 2.4 GHz, which makes it useful for dense networking but more difficult and costly to deploy outside of the home. Since the technology advancements described above require spectrum as a fundamental enabler, we must also explore how spectrum access may keep pace with technology.

Several changes to Wi-Fi spectrum access are possible to enable adequate capacity, coverage, and reliability necessary for wireless video delivery. Expanded and harmonized access to 5 GHz will enable significant capacity and full functionality of 802.11ac. Expanded bandwidth will also aid reliability, and additional capacity with unique interference protections may become available in the 3.5 GHz band. Finally, Wi-Fi coverage may be enhanced through new access to low-band unlicensed spectrum as part of the FCC’s broadcast incentive auction initiative. However, various challenges exist to realizing expanded spectrum access in each of these contexts.

5 GHz

As directed by Congress, the FCC is contemplating expanding Wi-Fi access in 5 GHz where it is not currently permitted, and is studying the effects of relaxing regulatory limits to Wi-Fi access in currently-available 5 GHz spectrum.

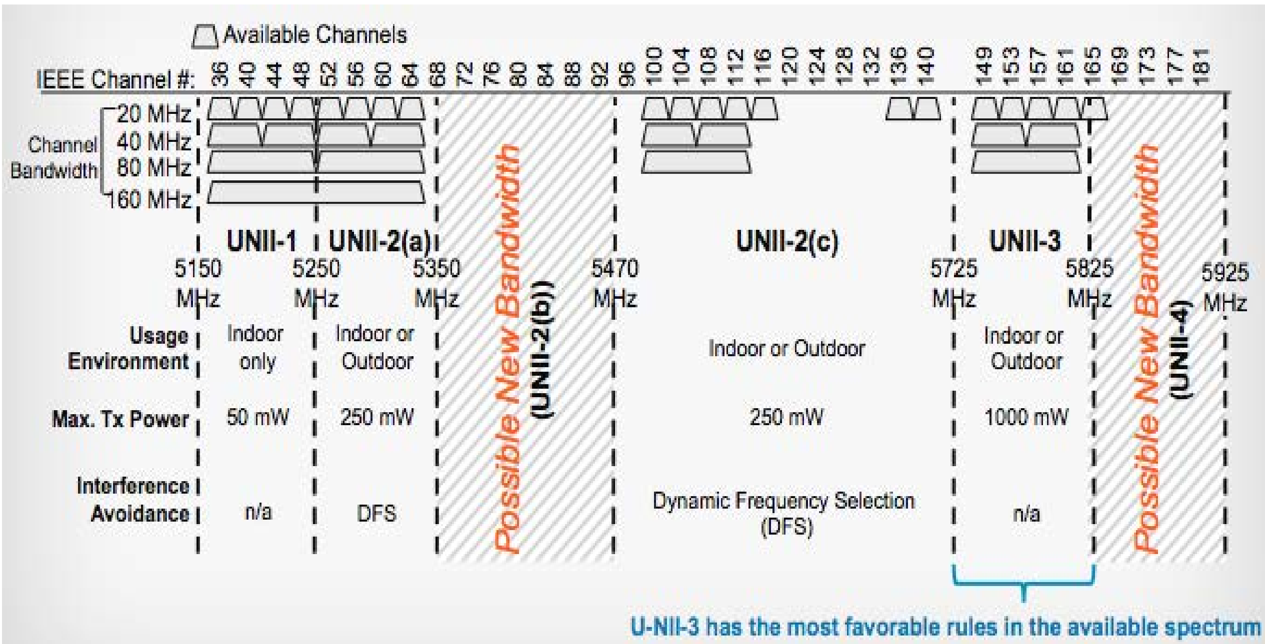

Figure 3: 5 GHz Wi-Fi Band Plan and Associated Rules

The 5 GHz U-NII band is split between an upper and lower portion, and different regulatory limits are placed on different portions of the band. In general, the lower channels are more restrictive in terms of power limits and technology, in order to avoid interference with radar and other incumbent systems that share the band. Only one portion of the band, 100 MHz known as “UNII-3”, has terms of access similar to 2.4 GHz, and most utilization of the band to date has centered around that portion. Figure 3 shows specific regulatory differences across the 5 GHz band. As shown, full implementation of 802.11ac using a 160 MHz channel is not possible in the United States under the current rules without crossing sub-bands that have different rules. This limits the utility of the band and of new Wi-Fi standards.

However, expanding 5 GHz Wi-Fi bandwidth and harmonizing access rules would enable the significant technology advances explored in this paper. In particular, expanding access above 5.9 GHz has the potential to enable an additional 100 MHz of flexible Wi-Fi spectrum adjacent to the "UNII-3" band, providing 200 MHz of contiguous Wi-Fi spectrum with terms of access that are similar to 2.4 GHz.

A similar approach is also possible in amending the terms of access to "UNII-1". "UNII-2" has to date been treated differently from other Wi-Fi spectrum in its requirements for dynamic frequency selection (DFS) for interference avoidance vis-à-vis incumbent radar systems. However, DFS is not widely supported by vendors because of this unique requirement and its uncertain performance.

In addition, currently-allocated Wi-Fi spectrum in "UNII-2c" has additional restrictions on use due to weather radar interference. Figure 4 below shows a possible new band plan for 5 GHz with expanded access to enable full utilization of 802.11ac and related technology.

Expanded, harmonized Wi-Fi access at 5 GHz would provide significant capacity and enable new technology. However, this outcome is only possible if appropriate, commercially-flexible coexistence protocols can be determined with incumbent spectrum users. Figure 4 below shows a possible new band plan for 5 GHz with expanded access, as well as incumbent spectrum allocations across the 5 GHz band. The number and diversity of incumbents outlines the challenge to Wi-Fi use of the 5 GHz band.

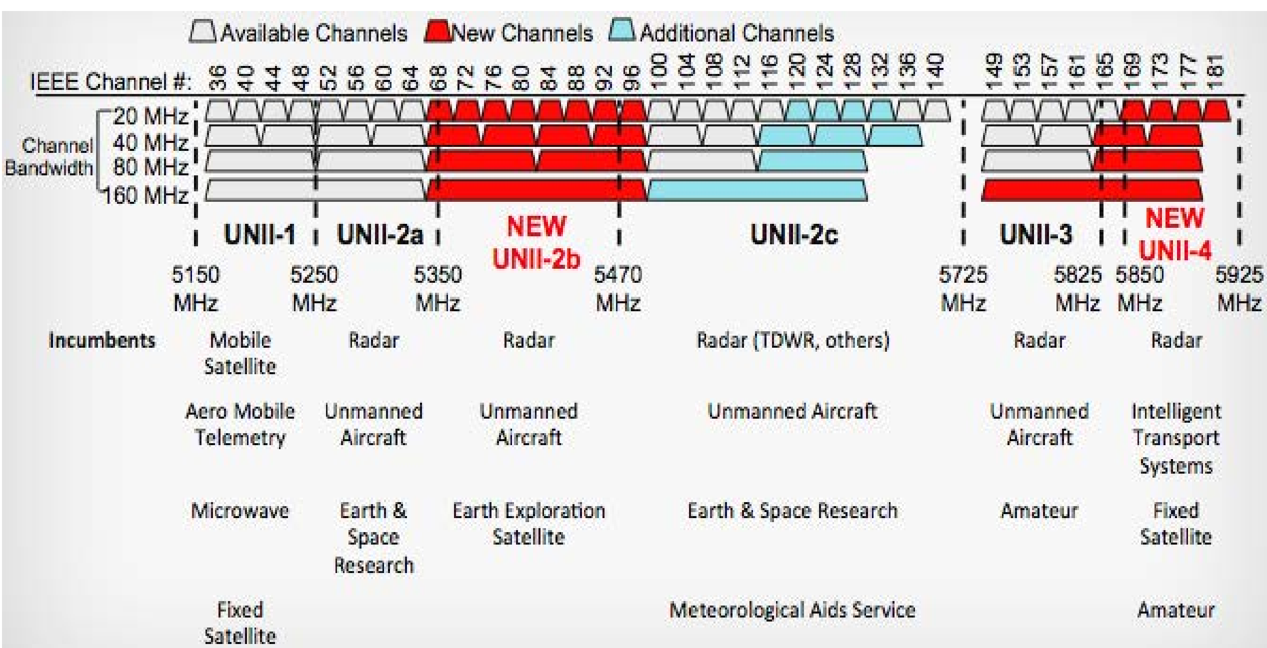

Figure 4: Possible New 5 GHz Wi-Fi Band and Non-Wi-Fi Spectrum Incumbents

3.5 GHz

In November 2010, the Department of Commerce recommended that the 3.55-3.65 GHz band be made available for new commercial use, and in December 2012, the FCC released a proposal for doing just that. Under the FCC's proposal, 100 MHz of spectrum would be made available under a 'general access' approach, akin to unlicensed, and possibly under a 'priority access' protocol with certain interference protections. This band may serve therefore to supplement Wi-Fi capacity and, if made available under priority access terms, could also provide reliability to Wi-Fi systems.

In recommending that the 3.5 GHz band be made commercially available, the Commerce Department noted that such availability would likely entail geographic restrictions to protect incumbent military radars that use the band. There are no plans to move these radars to another band due to prohibitive cost. Therefore, new commercial systems will be required to work around them, and operate at sufferance. In combination with its propagation characteristics, these circumstances limit the interest of mobile operators in utilizing the band, which provides an opportunity for economic terms of access under a "small cell" framework that may complement the Wi-Fi ecosystem.

The government proposed to manage commercial – military sharing in 3.5 GHz through broad "exclusion zones" that would preclude commercial use in the excluded area. These exclusions are shown in Figure 5 below, and are estimated to exclude approximately 60% of the US population from coverage.

These exclusion zones were developed through the government's worst-case interference analysis, assuming that new commercial services would be high-powered macro mobile networks. However, under a small cell deployment, there is reason to believe that the zones could be reduced substantially. In addition, with innovative access protocols, such as dynamic power adaptation and temporal coordination of signaling, commercial exclusion zones may be eliminated.

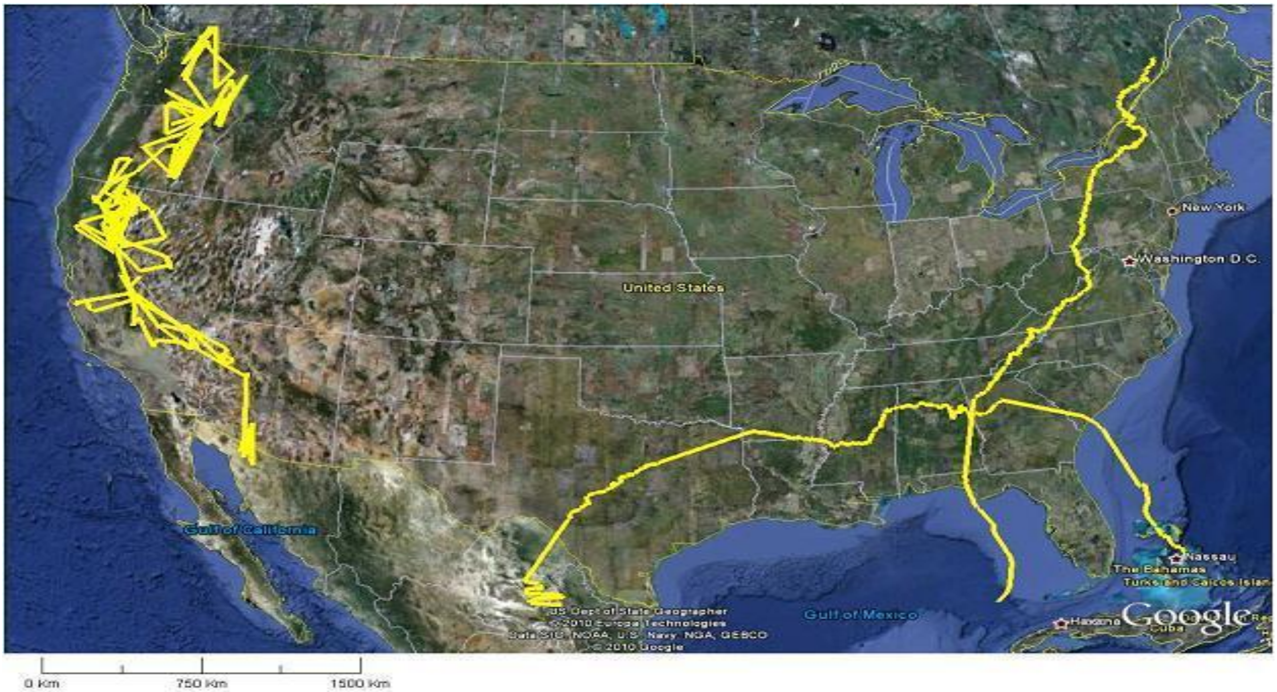

Figure 5: 3.5 GHz Commercial Exclusion from Costal Zones As Proposed by the Commerce Department [xiv]

Ongoing engagement with the Commerce Department and the FCC will be required to enhance the prospects of Wi-Fi use of the 3.5 GHz band.

600 MHz

Under Congressional direction, the FCC is preparing to reallocate up to 120 MHz of the broadcast television band for new wireless broadband use through a process known as an incentive auction. This process is, in essence, a two-sided auction in which the FCC buys spectrum back from broadcasters, repackages it for wireless use (including developing a new channelization plan for remaining broadcasters), and sells it.

In developing a new band plan for the repurposed spectrum, guard bands will be needed to protect mobile and broadcast uses. Unlicensed use is envisioned for these guard bands. If achieved, unlicensed use of 600 MHz guard bands would provide new geographic contiguity in low-band unlicensed spectrum, which is missing from the white space ecosystem today, hindering its development and utility.

Low-band unlicensed spectrum may prove to be an essential element of the Wi-Fi future. As access points are deployed outside of the home to achieve broader coverage, the favorable propagation and path characteristics of the 600 MHz band will enable a wider footprint on a more economic basis than spectrum higher in frequency. All else being equal, the 600 MHz band propagates over 60 times farther than 5 GHz, and over 14 times farther than 2.4 GHz. This property greatly reduces the number of access points needed to achieve coverage. For example, to cover an area equivalent to a 1W 600 MHz access point would require approximately 2900 2.4 GHz access points.[xv]

Such favorable propagation characteristics and related deployment economics would greatly speed Wi-Fi coverage, complementing many of the technology advances on the horizon, and enabling greater utility for a range of applications, including video. In addition, the new availability of low-band unlicensed spectrum, contiguous in frequency and geography, could enable simplification of access protocols that burden the 'White Spaces' today. The use of a 'control channel' architecture in low-band spectrum could enable network signaling to Wi-Fi devices across the full geographic footprint, both in 600 MHz and in other bands, that could greatly increase the availability and utility of unlicensed Wi-Fi spectrum.

However, the extent of unlicensed access in 600 MHz will be determined by the amount of spectrum supplied by broadcasters in the incentive auction, and the band plan the FCC chooses for the repurposed spectrum. Spectrum supply will be determined by decisions of broadcasters in the auction, but the band plan framework will be known in advance. The FCC has proposed a band plan framework that would yield very little guard band unlicensed spectrum – as little as 12 MHz, broken into two 6 MHz blocks, as seen in Figure 6 below.

Figure 6: FCC 600 MHz Band Plan Proposal[xvi]

An alternative band plan approach would be more likely to yield useful unlicensed spectrum while also better protecting mobile and TV services from interference. Such an approach would entail clearing spectrum from the top of the TV band and working down as spectrum is supplied by broadcasters, with the mobile uplink and downlink separated by an unlicensed duplex gap. This approach is depicted in Figure 7 below, which the FCC calls "Down from 51".

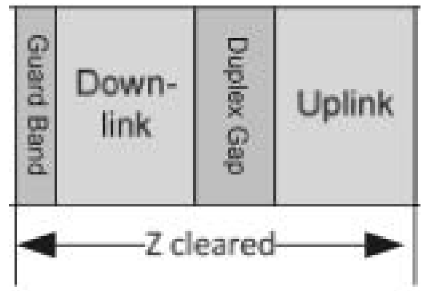

Figure 7: FCC Alternative Band Plan, "Down from 51"[xvii]

An appropriately sized duplex gap could yield 24 MHz of unlicensed spectrum, as noted by Grunwald and Baker (2013) and shown in Figure 8.[xviii]

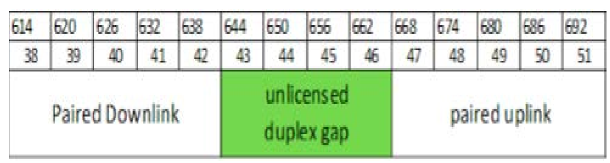

Figure 8: "Down from 51" Band Plan Yielding 24 MHz of Unlicensed Access[xix]

While there appears to be broad industry agreement on the merits of a "Down from 51" approach, the precise amount of unlicensed spectrum that would be appropriate remains under debate. The outcome of this debate will be important to realizing the benefits of low-band unlicensed spectrum and the new Wi-Fi technologies that consumers will value.

Spectrum Roadmap

The three wireless bands noted here – 600 MHz, 3.5 GHz, and 5 GHz – serve as a significant opportunity for evolving the Wi-Fi experience, and mitigating the risk of reduced performance as congestion grows in existing Wi-Fi bands.

With newfound coverage-related spectrum in 600 MHz, the ubiquity of Wi-Fi will be enhanced. With basic interference protections in the 3.5 GHz band, greater reliability may become possible for users. And with expanded and simplified access in 5 GHz, significant capacity will be achieved. These advances complement the technology developments on the Wi-Fi horizon.

Realizing Wi-Fi access to these bands presents unique challenges in each context, and will require sustained engagement with regulators and other stakeholders to achieve success.

PUTTING IT TOGETHER FOR THE CONSUMER

This paper has explored a range of next-generation Wi-Fi technologies and new spectrum access frameworks that may enable them. We have also highlighted some of the key dependencies for realizing this future vision of Wi-Fi.

To explore how these developments will be useful to consumers, we can examine how a hypothetical consumer may use video applications over Wi-Fi in the course of a day. In so doing, it becomes clear that our hypothetical user – and Wi-Fi users everywhere – will enjoy more consistent and reliable throughput and coverage, as well as more seamless network connection. These developments will lead consumers to use Wi-Fi in a wider range of contexts, helping to ensure that Wi-Fi remains the broadband access network of choice.

Our use case progression tracks our consumer through a typical morning as he uses video over Wi-Fi, from the time he wakes up to his morning meeting in the office. The stages of our use case progression are described below, along with a brief description of the operating environment and the key evolutions in Wi-Fi technology and spectrum that will improve his experience.

Figure 9: Use Case Progression: A Morning of Video Over Wi-Fi

Stage 1, "Waking up": Our user begins his day by watching the news on his iPad as he gets out of bed.		
Operating Environment: Indoor home use; access point dedicated to single household.		
Current Experience	*Future Experience*	*Enablers*
Manual setup, 'Best effort' quality	Automatic setup, Consistent high quality	802.11ac: wide channels, MIMO, SON Additional 5 GHz spectrum

Stage 2, "Eating Breakfast": Our user moves to the living room to eat in front of the TV.		
Operating Environment: Indoor home use; access point dedicated to single household.		
Current Experience	*Future Experience*	*Enablers*
Uses set-top to locate and view linear content	Mirrors iPad content on TV wirelessly	Miracast DLNA Gateway WiGig Additional 5 GHz spectrum

Stage 3, "Out the door": Our user walks to the train station, listening to the audio of video content continuing to stream on mobile device or iPad.		
Operating Environment: Outdoor neighborhood use; shared access point on the street.		
Current Experience	*Future Experience*	*Enablers*
Manual reconnection to APs, spotty coverage and throughput. Possible default to mobile.	Automatic AP connection, more consistent experience. Consistent experience enables continued Wi-Fi use.	802.11ac Passpoint / 802.11r SON / 802.11k New spectrum bands and standards (3.5 GHz for reliability, 600 MHz for coverage)

Stage 4, "On the train": Our user continues to view video content on his train ride to work.		
Operating Environment: Shared use inside train car with mobility; antenna on car connects to rail wayside access point.		
Current Experience	*Future Experience*	*Enablers*
Manual connection, irregular connectivity due to mobility, possible congestion due to	Automatic connection, consistent experience despite mobility, shared AP. Seamless experience enables	802.11ac Passpoint / 802.11r SON / 802.11k New spectrum bands and standards (3.5 GHz for reliability, 600 MHz for coverage and mobility)

shared use. Likely default to mobile use.	continued use of Wi-Fi.	

Stage 5, "Walking to work": Our user walks from train station to office, listening to the audio of video content continuing to stream on mobile device or iPad.		
Operating Environment: Outdoor downtown use; access point shared with many commuters.		
Current Experience	*Future Experience*	*Enablers*
Manual connection to SMB or public APs, low or no throughput due to congestion. Possible default to mobile.	Automatic connection to available APs, connectivity maintained despite congestion. Lower risk of default to mobile.	802.11ac Passpoint / 802.11r SON / 802.11k New spectrum bands and standards (3.5 GHz for reliability, 600 MHz for coverage)

Stage 6, "Morning meeting": Our user arrives at the office and uses Wi-Fi to facilitate a client meeting, delivering a presentation.		
Operating Environment: Indoor access point dedicated to meeting room or floor.		
Current Experience	*Future Experience*	*Enablers*
Uses Wi-Fi to pull presentation to laptop, then manually connects to projector, and emails presentation to clients later.	Uses Wi-Fi to pull presentation to laptop, WiGig to connect wirelessly to monitor, and Miracast to enable clients to view on their own screens in real time.	802.11ac Passpoint / 802.11r SON / 802.11k WiGig Miracast New spectrum bands and standards (3.5 GHz for reliability, 600 MHz for coverage)

CONCLUSION

The typical Wi-Fi user will benefit from a range of new technologies and new spectrum, enhancing capacity, reliability, and coverage, and making Wi-Fi more useful in a number of contexts. In our examination of one use case progression, these factors enable Wi-Fi to serve new purposes for consumers. Through adoption of new standards and optimization techniques, Wi-Fi can enable new home uses, and can become a more consistent and seamless experience for users outside of the home. Key to this evolution is adoption of new technologies in consumer devices and new spectrum access to enable these advances. Sustained engagement in the Wi-Fi ecosystem and spectrum policy on the part of MSOs and the broader Wi-Fi community is therefore of critical importance in realizing the future of Wi-Fi.

Endnotes

[i] Sandvine, "Global Internet Phenomena Report", 2H2012. The traffic category of "real time entertainment" comprises 65.2% of peak-period fixed access traffic in North America. Sandvine notes that this traffic category has, overall, doubled its share of total traffic over the past three years.

[ii] Cisco forecasts North American Internet traffic to grow 2.8 times between 2011 and 2016, while peak-period traffic is estimated to grow 3.1 times. Cisco Visual Networking Index, 2011-2016.

[iii] Business Insider, "How Tablet Sales Have Exploded, And Where The Next Wave of Growth Will Come From", January 8, 2013;.

[iv] Gartner Research, 2012, as reported by Matt Hambien, "Smartphones and tablets growth exploding, especially in business, Gartner says", Computerworld, November 6, 2012.

[v] Chetan Sharma Consulting, "US Wireless Market Update, Q42011 and Full Year 2011". Note that mobile-enabled tablets generally also come equipped with Wi-Fi.

[vi] The terms MSO and Cable operator are used interchangeably in this paper.

[vii] NCTA, "In America, There Are 10x More Cable Operator Provided Wi-Fi Hotspots Than There Are Starbucks", March 28, 2013.

[viii] Previous efforts at municipal Wi-Fi networks existed but largely failed due to an inability to support the investment required to build and operate the networks. For a review of municipal Wi-Fi, see Tim Wu, "Where's My Free Wi-Fi?", Slate, September 27, 2007.

[ix] Sandvine, 2012.

[x] Broadcom, "IEEE 802.11ac – Wi-Fi for the Mobile and Video Generation", January 2012.

[xi] For more information, see Wi-Fi Alliance, "Wi-Fi Alliance and Wireless Gigabit Alliance to Unify", press release, January 3, 2013

[xii] ABI Research, "Increasing DLNA Software Certification Will Propel the Adoption and Connection of Devices within the Home Network", January 24, 2011.

[xiii] Wi-Fi Multimedia is a WFA certification based on the IEEE 802.11e standard, which enables traffic prioritization by category.

[xiv] Department of Commerce, National Telecommunications and Information Administration, "An Assessment of the Near-Term Viability of Accommodating Wireless Broadband Systems in the 1675-1710 MHz, 1755-1780 MHz, 3500-3650 MHz, and 4200-4220 MHz, 4380-4400 MHz Bands", October 2010.

[xv] Dirk Grunwald and Kenneth Baker, "FCC Broadcast Incentive Auction: A Band Plan Framework for Maximizing Spectrum Utility", 2013.

[xvi] FCC 12-118, "In the Matter of Expanding the Economic and Innovation Opportunities of Spectrum Through Incentive Auctions", Notice of Proposed Rulemaking, released October 2, 2012.

[xvii] Ibid.

[xviii] Grunwald and Baker, 2013.

[xix] In this band plan and the others shown in this document, only a portion of the band is depicted due to space limitations. In the Grunwald and Baker example, 120 megahertz

of spectrum is assumed to be repurposed through the incentive auction; additional mobile spectrum is assumed in the band plan but not shown here.

QoE: As Easy As PIE

Rong Pan, Preethi Natarajan, Chiara Piglione, Mythili S. Prabhu, Alon Bernstein
Advanced Architecture & Research Group,
Cisco Systems Inc., San Jose, CA 95134, U.S.A.

***Abstract*—Bufferbloat is a phenomenon where excess buffers in the network, such as in cable modems or wireless APs, cause high latency and jitter. As more and more interactive applications (e.g. voice over IP and real time video conferencing) run in the Internet, high latency and jitter degrade application performance and impact users' Quality of Experience (QoE). There is a pressing need to design intelligent queue management schemes that can control latency and jitter; and hence provide desirable quality of service to users.**

We present here a lightweight design, PIE (Proportional Integral controller Enhanced), that can effectively control the average queueing latency to a reference value. The design does not require per-packet extra processing, so it incurs very small overhead and is simple to implement in both hardware and software. In addition, the design parameters are self-tuning, and hence PIE is robust and optimized for various network scenarios. We apply the algorithm in DOCSIS 3.0 environment. Simulation results show that PIE can ensure low latency and achieve high link utilization under various congestion situations.

***Index Terms*—bufferbloat, Active Queue Management (AQM), Quality of Service (QoS), Quality of Experience (QoE), Explicit Congestion Notification (ECN)**

I. Introduction

The explosion of smart phones, tablets and video traffic in the Internet brings about a unique set of challenges for congestion control. To avoid packet drops, many service providers or data center operators require vendors to put in as much buffer as possible. With rapid decrease in memory chip prices, these requests are easily accommodated to keep customers happy. However, the above solution of large buffers fails to take into account the nature of TCP, the dominant transport protocol running in the Internet. The TCP protocol continuously increases its sending rate and causes network buffers to fill up. TCP cuts its rate only when it receives a packet drop or mark that is interpreted as a congestion signal. However, drops and marks usually occur when network buffers are full or almost full. As a result, excess buffers, initially designed to avoid packet drops, would lead to highly elevated queueing latency and jitter. The phenomenon was detailed in 2009 [1] and the term, "bufferbloat" was introduced by Jim Gettys in late 2010 [2]. Recent studies of home access network also confirmed that modems often have large buffers and that DSL links often have large high latency [3], [4].

Active queue management (AQM) schemes, such as RED [5], BLUE [6], PI [7], AVQ [8], etc, have been around for well over a decade. By selectively dropping packets early to optimize traffic behaviors, AQM schemes could potentially solve the aforementioned problem. RFC 2309 [9] strongly recommends the adoption of AQM schemes in the network to improve the performance of the Internet. Although controlling delay is a key metric in assuring QoS, the DOCSIS specifications do not define the adoption of AQM and only recently added the option to set a small queue size. The rationale behind it was that DOCSIS does not need AQM because DOCSIS addresses head-of-line blocking with fine-grained queuing or scheduling, i.e. sorting traffic into individual queue and then scheduling the drain of the queues according to the priority of these queues. In this way, the delay of the high priority traffic is minimized since it does not incur the delay issues caused by head-of-line blocking.

There are a few issues with the above line of thinking. First, head-of-line blocking is not the only cause for delay. Even if there is no head-of-line blocking, a large queue can still cause a large delay. For example, voice traffic can be sorted to one queue and data traffic can be put into a separate queue so there is no blocking of voice traffic. Nonetheless the data queue can still experience excessive delay. Second, limiting the buffer size is not the solution to the bufferbloat problem because latency is still high under persistent congestion. In addition, small buffer does not have enough space to absorb short-term bursts, which could lead to throughput loss. Third, even with today's technology, per flow queueing is still rather complicated to implement. Except for traffic class like voice, a lot of flows are queued into to the same queue. For example, a background e-mail download can still create head-of-line blocking for a web browsing session. Latency control for this data queue is still required in order to guarantee QoE. Due to aforementioned issues, we believe that AQM is crucial in controlling latency and solving the bufferbloat problem.

Although AQM has not been adopted in the DOCSIS specifications, RED, a well-known AQM scheme, has been adopted in a wide variety of network devices, such as switches and routers; and it has been implemented in both hardware and software. Unfortunately, due to the fact that RED needs careful tuning of its parameters for various network conditions, most network operators do not turn RED on. In addition, RED is designed to control the queue length which would affect delay implicitly. It does not control latency directly.

We recognize that the delay bloat caused by poorly managed big buffers is the core issue here. If latency can be controlled, bufferbloat, i.e., adding more buffers for bursts, is not a problem. More buffer space would allow larger bursts of packets to pass through as long as we control the average queueing delay to be small. Unfortunately, Internet today still lacks an effective design that can control buffer latency to improve QoE of latency-sensitive applications. In addition, it is a delicate balancing act to design a queue management scheme that not only allows short-term burst to smoothly pass, but also controls the average latency when long-term congestion persists.

Recently, a new AQM scheme, CoDel [10], was proposed to control the latency directly to address the bufferbloat problem. CoDel requires per packet timestamps. Also, packets are dropped at the dequeue function after they have been enqueued for a while. Both of these requirements consume excessive processing and infrastructure resources. This consumption will make CoDel expensive to implement and operate, especially in hardware.

In this paper, we present a lightweight algorithm, PIE (Proportional Integral controller Enhanced), which combines the benefits of both RED and CoDel: easy to implement like RED while directly control latency like CoDel. Similar to RED, PIE randomly drops a packet at the onset of the congestion. The congestion detection, however, is based on the queueing latency like CoDel instead of the queue length like conventional AQM schemes such as RED. Furthermore, PIE also uses the latency moving trends: latency increasing or decreasing, to help determine congestion levels. In addition, the design parameters are self-tuning, and hence PIE is robust and optimized for various network scenarios. Our DOCSIS 3.0 simulation results show that PIE can control latency around the reference under various congestion conditions. Furthermore, it can quickly and automatically respond to network congestion changes in an agile manner.

In what follows, Section II specifies our goals of designing the latency-based AQM scheme. Section III explains the scheme in detail. Section IV presents simulation results of the proposed scheme. In Section V, we discuss the implementation cost of PIE. Section VII concludes the paper and discusses future work.

II. Design Goals

We explore a queue management framework where we aim to improve the performance of interactive and delay-sensitive applications. The design of our scheme follows a few basic criteria.

- *Low Latency Control.* We directly control queueing latency instead of controlling queue length. Queue sizes change with queue draining rates and various flows' round trip times. Delay bloat is the real issue that we need to address as it impairs real time applications. If latency can be controlled to be small, bufferbloat is not an issue. In fact, we would allow more buffers for sporadic bursts as long as the average latency is under control.

- *High Link Utilization.* We aim to achieve high link utilization. The goal of low latency shall be achieved without suffering link under-utilization or losing network efficiency. An early congestion signal could cause TCP to back off and avoid queue buildup. On the other hand, however, TCP's rate reduction could result in link under-utilization. There is a delicate balance between achieving high link utilization and low latency.

- *Simple Implementation.* The scheme should be simple to implement and easily scalable in both hardware and software. The wide adoption of RED over a variety of network devices is a testament to the power of simple random early dropping/marking. We strive to maintain similar design simplicity.

- *Guaranteed Stability and Fast Responsiveness.* The scheme should ensure system stability for various network topologies and scale well with arbitrary number streams. The system also should be agile to sudden changes in network conditions. Design parameters shall be set automatically. One only needs to set performance-related parameters such as target queue delay, but no need to set any of the design parameters.

We aim to find an algorithm that achieves the above goals. It is noted that, although important, fairness is orthogonal to the AQM design whose primary goal is to control latency for a given queue. Techniques such as Fair Queueing [11] or its approximate such as Stochastic Fair Queueing (SFQ) [12] can be combined with any AQM scheme to achieve fairness. Therefore, in this paper, we focus on controlling a queue's latency and ensuring flows' fairness is not worse than those under the standard DropTail or RED design.

III. The PIE Scheme

In the section, we describe in detail the design of PIE and its operations. As illustrated in Figure 1, our scheme comprises three simple components: a) random dropping at enqueueing; b) periodic drop probability update; c) departure rate estimation.

The following subsections describe these components in further details, and explain how they interact with each other. At the end of this section, we will discuss how the scheme can be easily augmented to precisely control bursts.

A. Random Dropping

Like most state-of-the-art AQM schemes, PIE would drop packets randomly according to a drop probability, p, that is obtained from the "drop probability calculation" component. No extra step, like timestamp insertion, is needed. The procedure is as follows:

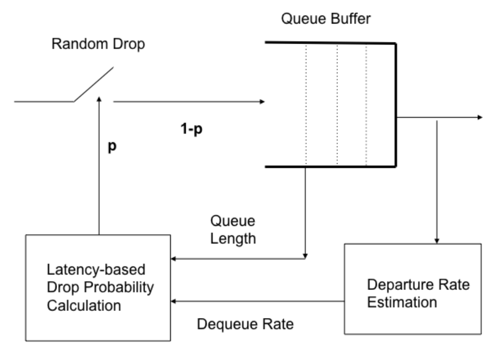

Fig. 1. Overview of the PIE Design. The scheme comprises three simple components: a) random dropping at enqueueing; b) latency based drop probability update; c) departure rate estimation.

Random Dropping:

Upon packet arrival

randomly drop a packet with a probability p.

B. Drop Probability Calculation

The PIE algorithm updates the drop probability periodically as follows:

- estimate current queueing delay using Little's law:
  $cur\_del = \frac{qlen}{avg\_drate}$;

- calculate drop probability p as:
  $p = p + \alpha * (cur\_del - ref\_del) + \beta * (cur\_del - old\_del)$;

- update previous delay sample as:
  $old\_del = cur\_del$.

The average draining rate of the queue, avg_drate, is obtained from the "departure rate estimation" block. Variables, cur_del and old_del, represent the current and previous estimation of the queueing delay. The reference latency value is expressed in ref_del. The update interval is denoted as T_{update}. Parameters α and β are scaling factors.

Note that the calculation of drop probability is based not only on the current estimation of the queueing delay, but also on the direction where the delay is moving, i.e., whether the delay is getting longer or shorter. This direction can simply be measured as the difference between cur_del and old_del. Parameter α determines how the deviation of current latency from the target value affects the drop probability; β exerts the amount of additional adjustments depending on whether the latency is trending up or down. The drop probability would be stabilized when the latency is stable, i.e. cur_del equals old_del; and the value of the latency is equal to ref_del. The relative weight between α and β determines the final balance between latency offset and latency jitter. This is the classic Proportional Integral controller design [13], which has been adopted for controlling the queue length before in [7] and [14]. We adopt it here for controlling queueing latency. In addition, to further enhance the performance, we improve the design by making it auto-tuning as follows:

if $p < 1\%$: $\alpha = \tilde{\alpha}/8$; $\beta = \tilde{\beta}/8$;
else if $p < 10\%$: $\alpha = \tilde{\alpha}/2$; $\beta = \tilde{\beta}/2$;
else: $\alpha = \tilde{\alpha}$; $\beta = \tilde{\beta}$;

where $\tilde{\alpha}$ and $\tilde{\beta}$ are static configured parameters. Auto-tuning would help us not only to maintain stability but also to respond fast to sudden changes. The intuitions are the following: to avoid big swings in adjustments which often leads to instability, we would like to tune p in small increments. Suppose that p is in the range of 1%, then we would want the value of α and β to be small enough, say 0.1%, adjustment in each step. If p is in the higher range, say above 10%, then the situation would warrant a higher single step tuning, for example 1%. The procedures of drop probability calculation can be summarized as follows.

Drop Probability Calculation:

Every T_{update} interval

1. Estimation current queueing delay:

$$cur\_del = \frac{qlen}{avg\_drate}.$$

2. Based on current drop probability, p, determine suitable step scales:

$$\begin{aligned} if\ p\ <\ 1\%, \quad \alpha = \ \tilde{\alpha}/8; \quad \beta = \tilde{\beta}/8; \\ else\ if\ p\ < 10\%, \quad \alpha = \ \tilde{\alpha}/2; \quad \beta = \tilde{\beta}/2; \\ else\ , \quad \alpha = \ \tilde{\alpha}; \quad \beta = \tilde{\beta}; \end{aligned}$$

3. Calculate drop probability as:

$$\begin{aligned} p = p + \alpha \ * \ (cur\_del - ref\_del) + \\ \beta \ * \ (cur\_del - old\_del); \end{aligned}$$

4. Update previous delay sample as:

$$old\_del = cur\_del.$$

We have discussed packet drops so far. The algorithm can be easily applied to networks codes where Early Congestion Notification (ECN) is enabled. The drop probability p could simply mean marking probability.

C. Departure Rate Estimation

The draining rate of a queue in the network often varies either because other queues are sharing the same link, or the link capacity fluctuates. Rate fluctuation is particularly common in wireless networks. Hence, we decide to measure the departure rate directly as follows:

Departure Rate Calculation:

Upon packet departure

1. Decide to be in a measurement cycle if:

$$qlen > dq\_threshold;$$

2. If the above is true, update departure count dq_count:

$$dq\_count = dq\_count + dq\_pktsize;$$

3. Update departure rate once $dq\_count > dq\_threshold$ and reset counters:

$$\begin{aligned} dq\_int &= now - start; \\ dq\_rate &= \frac{dq\_count}{dq\_int}; \\ avg\_drate &= (1-\varepsilon) * avg\_drate \\ &\quad + \ \varepsilon * dq\_rate \\ start &= now. \\ dq\_count &= 0; \end{aligned}$$

From time to time, short, non-persistent bursts of packets result in empty queues, this would make the measurement less accurate. Hence we only measure the departure rate, dq_rate, when there are sufficient data in the buffer, i.e., when the queue length is over a certain threshold, $dq\_threshold$. Once this threshold is crossed, we obtain a measurement sample. The samples are exponentially averaged, with averaging parameter ε, to obtain the average dequeue rate, avg_drate. The parameter, dq_count, represents the number of bytes departed since the last measurement. The threshold is recommended to be set to 5KB assuming a typical packet size of around 1KB or 1.5KB. This threshold would allow us a long enough period, dq_int, to obtain an average draining rate but also fast enough to reflect sudden changes in the draining rate. Note that this threshold is not crucial for the system's stability.

D. Handling Bursts

The above three components form the basis of the PIE algorithm. Although we aim to control the average latency of a congested queue, the scheme should allow short term bursts to pass through the system without hurting them. We would like to discuss how PIE manages bursts in this section.

Bursts are well tolerated in the basic scheme for the following reasons: first, the drop probability is updated periodically. Any short term burst that occurs within this period could pass through without incurring extra drops as it would not trigger a new drop probability calculation. Secondly, PIE's drop probability calculation is done incrementally. A single update would only lead to a small incremental change in the probability. So if it happens that a burst does occur at the exact instant that the probability is being calculated, the incremental nature of the calculation would ensure its impact is kept small.

Nonetheless, we would like to give users a precise control of the burst. We introduce a parameter, max_burst, that is similar to the burst tolerance in the token bucket design. By default, the parameter is set to be 100ms. Users can certainly modify it according to their application scenarios. The burst allowance is added into the basic PIE design as follows:

Burst Allowance Calculation:

Upon packet arrival

1. If $burst\_allow > 0$

$$enque\ packet\ bypassing\ random\ drop;$$

Upon dq_rate *update*

2. Update burst allowance:

$$burst\_allow = burst\_allow - dq\_int;$$

3. if $p = 0$; and both cur_del and old_del less than
$ref\_del/2$, reset $burst\_allow$,

$$burst\_allow = max\_burst;$$

The burst allowance, noted by $burst\_allow$, is initialized to max_burst. As long as $burst\_allow$ is above zero, an incoming packet will be enqueued bypassing the random drop process. Whenever dq_rate is updated, the value of $burst\_allow$ is decremented by the departure rate update period, dq_int. When the congestion goes away, defined by us as p equals to 0 and both the current and previous samples of estimated delay are less than $ref\_del/2$, we reset $burst\_allow$ to max_burst.

IV. Performance Evaluation

In this section we present our ns-2 [15] simulation results for the scenario of DOCSIS cable modems. We first demonstrate the basic functions of PIE using a few basic scenarios, and then compare PIE and CoDel performance using various scenarios. We focus our attention on the following performance metrics: instantaneous queueing delay, drop probability, TCP throughput and link utilization.

The CM (upstream) is modeled as a single queue that implements both rate shaping using a token bucket algorithm (with parameters Max Sustained Rate, Traffic Burst and Peak Traffic Rate), and the Request-Grant DOCSIS MAC. The CMTS (downstream) is implemented as a single rate shaping queue with token bucket. The upstream queue is modeled as a PIE (or other AQM scheme) queue, while the downstream queue is implemented using a DropTail model.

In our simulation, the upstream token bucket parameters are set to the following values: Max Sustained Rate = 5Mbps, Traffic Burst = 10MB and Peak Traffic Rate = 20Mbps. The downstream token bucket parameters are set as follows: Max Sustained Rate = 20Mbps, Traffic Burst = 20MB and Peak Traffic Rate = 50Mbps. The RTT is 100ms and unless otherwise stated the buffer size is 480ms . We use both TCP and UDP traffic for our evaluation. All TCP traffic sources are implemented as TCP New Reno with SACK running FTP applications. UDP traffic is implemented using Constant Bit Rate (CBR) sources. Both UDP and TCP packets are configured to have a fixed size of 1500B. Unless otherwise stated the PIE parameters are configured as follow: ref_del = 20ms, T_{update} = 15ms, α = 0.25Hz, $\beta = 2.5Hz$, $dq\_threshold$ = 4500B, max_burst = 50ms.

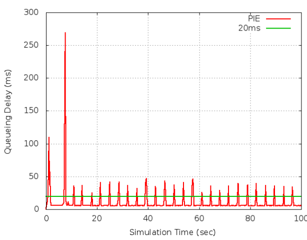

(a) Light 1 TCP Flows

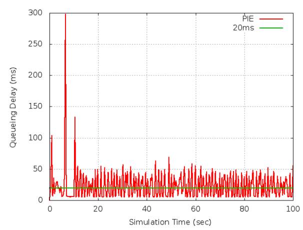

(b) Medium 5 TCP Flows

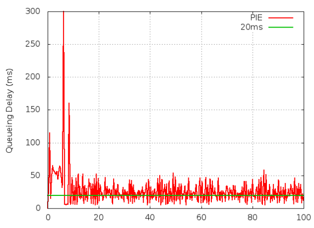

(c) Heavy 15 Flows

Fig. 2. Queueing Latency Under Various Traffic Loads: a) 1 TCP flow; b) 5 TCP flows; c) 15 TCP flows. Queueing latency is controlled at the reference level of 20ms regardless of the traffic intensity.

Function Verification: We first validate the functionalities of PIE, making sure it performs as designed using static traffic sources under various loads.

1) Light TCP traffic: We first consider a single TCP flow. Figure 2(a) and Figure 3(a) show the queueing delay and throughput, respectively. From Figure 3(a), it is clear that a single TCP flow is able to take advantage of the initial Peak Rate burst at 20Mbps except the initial dip around 1s. Once the token bucket transitions into the Maximum Sustained Rate of 5Mbps, we can see that the throughput oscillates around 5Mbps with the typical TCP sawtooth behavior. Due to dual token bucket rates, we are not losing throughput: if the throughput is under 5Mbps, the token of the peak rate would allow it to burst over 5Mbps. Due to the rate changes, there is a sudden spike in the queueing latency around 7s as shown in Figure 3(a). However,

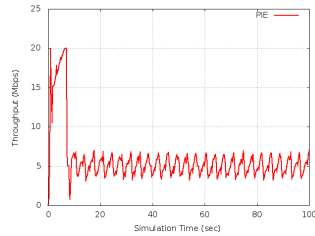

(a) Light 1 TCP Flows

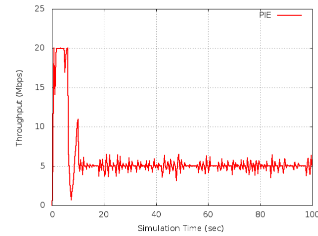

(b) Medium 5 TCP Flows

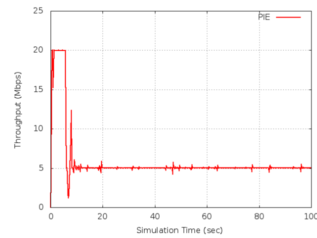

(c) Heavy 15 TCP Flows

Fig. 3. Link Throughput Under Various Traffic Loads: a) 1 TCP flow; b) 5 TCP flows; c) 15 flows. High link utilization is achieved regardless of traffic intensity, even under low multiplexing case.

the algorithm can quickly control the delay to be around the target value of 20ms. The average drop probability here is 0.9%.

2) Medium TCP traffic: In this test scenario, we increase the number of TCP flows to 5. With higher traffic intensity, the link utilization reaches 100% for both the peak and sustained rates as clearly shown in Figure 3(b). Except the sudden spike around 7s due to the rate change, the queueing delay, depicted in Figure 2(b), is controlled around the desired 20ms. The equilibrium latency is unaffected by the increased traffic intensity. Due to higher multiplexing, the link throughput is smoother. The queueing delay fluctuates more evenly around the reference level. The average throughput reaches full capacity of 5Mbps as shown in Figure 3(b). The average drop probability is 2.0%.

3) Heavy TCP traffic: To demonstrate PIE's performance under persistent heavy congestion, we increase the traffic load to 15 TCP flows. The corresponding latency plot can be found in Figure 2(c). Again, PIE is able to contain the queueing delay around the reference level regardless of the traffic mix while achieving both the peak and sustained rates shown in Figure 3(c). The average drop probability in this case is 6.2%.

4) PowerBoost: In this test, we investigate PIE's ability to control the queuing delay when the DOCSIS model shifts to a higher upstream bandwidth:

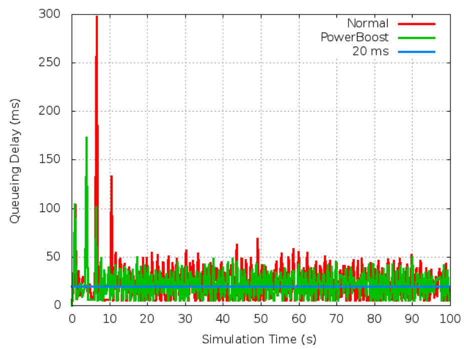

Fig. 4. Normal vs. Powerboost Delay Comparison Under 15 TCP flows: the queuing delay oscillating around the delay reference of 20ms in a similar fashion for both cases regardless of configured speeds.

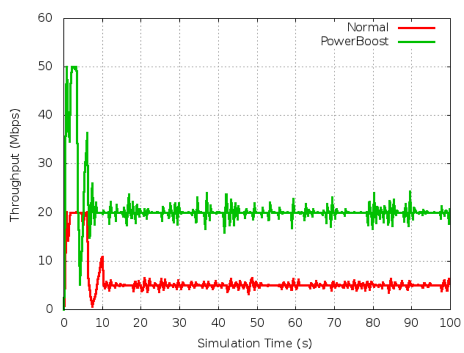

Fig. 5. Normal vs. Powerboost Throughput Comparison with 15 TCP flows: both the peak rate and the sustained rate are achieved under normal and Powerboost scenarios.

50Mbps as the peak rate and 20Mbps as the sustained rate (a.k.a PowerBoost). The critical aspect to verify is whether PIE's parameter settings hold for both the normal and PowerBoost scenarios.

Figure 4 and 5 plot the queuing delays experienced and throughputs achieved under the normal and PowerBoost conditions with 15 TCP flows. The plots show that, while different throughputs are achieved, the queuing delay oscillating around the delay reference of 20ms in a similar fashion for both cases. This verifies that PIE's auto-tuning helps PIE to adapt to the higher speed traffic conditions found in DOCSIS 3.0 cable modem environments.

Performance Evaluation and Comparison: The functions of PIE are verified above. This section evaluates PIE under various application scenarios, compares its performance against CoDel and shows

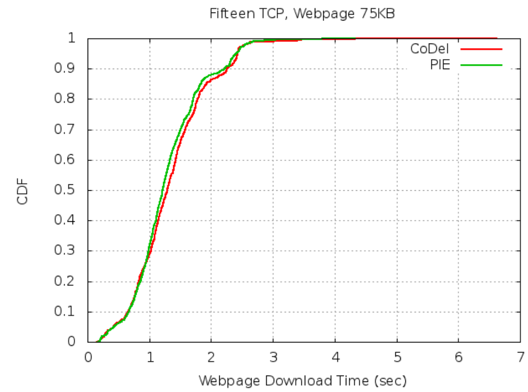

Fig. 6. PIE vs. Codel Performance Comparison: the CDF plot of 2000 web pages' download time with 15 long lived TCP flows running in the background. PIE and CoDel behave similarly in this situation.

how PIE is better suited for controlling latency in todays Internet. The simulation topology is similar to the above. The cable modem runs either the PIE or CoDel scheme.

5) HTTP Traffic: This test consists of 15 long-lived TCP flows that share the upstream bandwidth with 100 concurrent web page downloads. Each web page is 75KB in size and they are evenly downloaded from four different sites with RTTs of 20ms, 30ms, 50ms and 100ms respectively. The download process is repeated 20 times so that a total of 2000 web page downloads are generated. The long-lived TCP traffic has RTT of 100ms. Figure 6 shows the web page download time as CDF for both CoDel and PIE. From the graph, we see that the web page download times for both schemes are comparable.

6) Video traffic: This test consists of a single UDP flow at 6.5Mbps on the upstream link. Since most real time video communication adopt the UDP protocol, this test compares how each scheme behaves given a real-time, high-definition video traffic. Figure 7 shows the utilization of the upstream link for both CoDel and PIE. In both schemes, the link utilization is around 6.5Mbps (offered load) until 50s. Once tokens are exhausted at 50s, the link utilization goes down to 5Mbps for both schemes.

Figure 8 plots the queueing delay on the upstream queue. Both schemes only incur the MAC layer delay of 5ms-6ms in the initial 50s. Once the tokens are exhausted at 50s, queue builds up in both schemes. PIE is able to adapt to the increasing queue

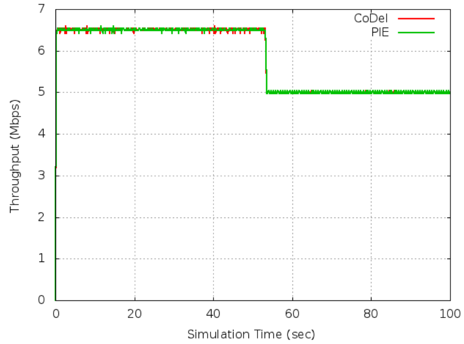

Fig. 7. PIE vs. Codel Performance Comparison Under UDP traffic: the test represents a real time video traffic which sends 6.5Mbps. In both schemes, the link utilization is around 6.5 Mbps (offered load) until 50s. Once the peak rate tokens are exhausted at 50s, the link utilization goes down to 5Mbps for both schemes.

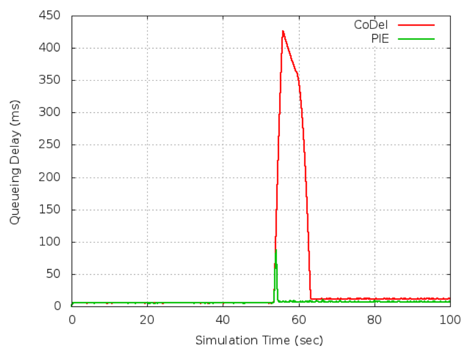

Fig. 8. PIE vs. Codel Performance Comparison Under UDP traffic: the queuing delay on the upstream queue. Both schemes maintain a low queuing delay in the initial 50s. Once the tokens are exhausted at 50s, queue builds up in both schemes. PIE is able to adapt to the increasing queue size and bring down the queuing latency around 53s, whereas CoDel takes much longer to bring down the latency (around 63s).

size and bring down the queuing latency around 53s, whereas CoDel takes much longer to bring down the latency (around 63s). PIEs auto-tuning features helps PIE to adapt faster to dynamically changing link and traffic conditions.

V. Implementation

PIE can be applied to existing hardware or software solutions. In this section, we discuss the implementation cost of the PIE algorithm. There are three steps involved in PIE as discussed in Section III. We examine their complexities as follows.

Upon packet arrival, the algorithm simply drops a packet randomly based on the drop probability p. This step is straightforward and requires no packet header examination and manipulation. Besides, since no per packet overhead, such as a timestamp, is required, there is no extra memory requirement. Furthermore, the input side of a queue is typically under software control while the output side of a queue is hardware based. Hence, a drop at enqueueing can be readily retrofitted into existing hardware or software implementations.

The drop probability calculation is done in the background and it occurs every T_{udpate} interval. Given modern high speed links, this period translates into once every tens, hundreds or even thousands of packets. Hence the calculation occurs at a much slower time scale than packet processing time, at least an order of magnitude slower. The calculation of drop probability involves multiplications using α and β. Since the algorithm is not sensitive to the values of α and β, we can choose the values, e.g. $\alpha = 0.25$ and $\beta = 2.5$ so that multiplications can be done using simple adds and shifts. As no complicated functions are required, PIE can be easily implemented in both hardware and software. The state requirement is only two variables per queue: cur_del and old_del. Hence the memory overhead is small.

In the departure rate estimation, PIE uses a counter to keep track of the number of bytes departed for the current interval. This counter is incremented per packet departure. Every T_{update}, PIE calculates latency using the departure rate, which can be implemented using a multiplication. Note that many network devices keep track an interface's departure rate. In this case, PIE might be able to reuse this information and incurs no extra cost. Besides, in cable modem or CMTS scenarios, the peak rate and sustained rate are preconfigured. Hence, PIE can take advantage of this rate information and current congestion state to simply skip the third step of the algorithm.

In summary, the state requirement for PIE is limited and computation overheads are small. Hence, PIE is simple to be implemented. In addition, since PIE does not require any user configuration, it does not impose any new cost on existing network management system solutions. SFQ can be combined with PIE to provide further improvement of latency for various flows with different priorities. However, SFQ requires extra queueing and scheduling structures. Whether the performance gain can justify the design overhead needs to be further investigated.

VI. Acknowledgement

We would like to thank Greg White from CableLabs for providing us the ns2 model for DOCSIS 3.0 cable modems.

VII. Conclusions and Future Work

In this paper we have described PIE, a latency-based design for controlling bufferbloat in the Internet. The PIE design bases its random dropping decisions not only on current queueing delay but also on the delay moving trend. In addition, the scheme self-tunes its parameters to optimize system performance. As a result, PIE is effective across diverse range of network scenarios. Our simulation studies of DOCSIS 3.0 modems show that PIE can ensure low latency under various congestion situations. It achieves high link utilization while maintaining stability consistently. It is a light-weight, enqueueing based design that works with both TCP and UDP traffic. The PIE design only requires low speed drop probability update, so it incurs very small overhead and is simple enough to implement in both hardware and software.

Going forward, we will explore efficient methods to provide weighted fairness under PIE. There are two ways to achieve this: either via differential dropping for flows sharing a same queue or through class-based fair queueing structure where flows are queued into different queues. There are pros and cons with either approach. We will study the tradeoffs between these two methods.

References

[1] B. Turner, "Has AT&T Wireless Data Congestion Been Self-Inflicted?" [Online]. Available: BroughTurnerBlog

[2] J. Gettys, "Bufferbloat: Dark buffers in the internet," *IEEE Internet Computing*, vol. 15, pp. 95–96, 2011.

[3] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxon, "Netalyzer: Illuminating the edge network," in *Proceedings of Internet Measurement Conference*, 2010.

[4] G. Maier, A. Feldmann, V. Paxon, and M. Allman, "On dominant characteristics of residential broadband internet traffic," in *Proceedings of Internet Measurement Conference*, 2009.

[5] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.

[6] W. Feng, K. Shin, D. Kandlur, and D. Saha, "The blue active queue management algorithms," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, pp. 513–528, Aug. 2002.

[7] C. V. Hollot, V. Misra, D. Towsley, and W. bo Gong, "On designing improved controllers for aqm routers support," in *Proceedings of IEEE Infocom*, 2001, pp. 1726–1734.

[8] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue (avq) algorithm for active queue management," in *Proceedings of ACM SIGCOMM*, 2001, pp. 123–134.

[9] B. Braden, D. Clark, J. Crowcroft, and et. al., "Recommendations on Queue Management and Congestion Avoidance in the Internet," RFC 2309 (Proposed Standard), 1998.

[10] K. Nichols and V. Jacobson, "A Modern AQM is just one piece of the solution to bufferbloat." [Online]. Available: http://queue.acm.org/detail.cfm?id=2209336

[11] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulaton of a fair queueing algorihtm," *Journal of Internetworking Research and Experience*, pp. 3–26, Oct. 1990.

[12] P. McKenney, "Stochastic fairness queueing," *Internetworking: Research and Experience*, vol. 2, pp. 113–131, Jan. 1991.

[13] G. Franklin, J. D. Powell, and A. Emami-Naeini, in *Feedback Control of Dynamic Systems*, 1995.

[14] R. Pan, B. Prabhakar, and et. al., "Data center bridging - congestion notification." [Online]. Available: http://www.ieee802.org/1/pages/802.1au.html

[15] "NS-2." [Online]. Available: http://www.isi.edu/nsnam/ns/

REMOTE STORAGE DVR

Stephanie Mitchko

Cablevision Systems LLC

Abstract

Approximately seven years ago the team at Cablevision envisioned a system where our subscribers' recorded content could be stored in a data center and played back over our network to their homes. The idea would eliminate the need for costly in home devices and move to a cloud environment. We built a prototype that used commodity hardware along with a combination of vendor and in house developed software that would allow subscribers to create recordings of their favorite shows and store those programs on storage systems in our headend. After announcing our intentions, the content owners filed suit against us, claiming that our system was not permissible under the copyright laws. After several years of litigation, the courts vindicated the system, holding that it did not violate the copyright laws.

In January of 2011 we commercially launched our Remote Storage DVR (RS-DVR) product and marketed it as DVR Plus. Region by region we enabled cloud based services for all of our customers. This paper reviews the overall system and addresses some lessons learned. This will include both technical and operational detail on how moving content recording and playback to the cloud enables new features and portability for our customers. It will also define how this platform is technically extensible to support multiple advanced streaming services as well as Dynamic Ad Insertion (DAI).

Contributors: *Thanks to Peter Caramanica, Rich Neil, Brad Feldman, and John Kenny for their contributions to this paper.*

There are four major areas that are addressed in this paper as part of the overall system,

- Product Features and Flexibility
- Client code architecture
- Software control plane
- Ingest Storage & Streaming

PRODUCT FEATURES AND FLEXIBILITY OF THE CLOUD

Tuners and Storage

The key concept in an RS-DVR solution is moving the storage and tuner functions out of the home. This eliminates the need for storage and tuners in Set-Top boxes and moves the features to the cloud. By removing the need for physical tuners and local storage we break away from the hardware-based restrictions of in home DVRs. Once these features are removed from the physical device in the home, the ability to configure flexible product offerings is technologically available.

For example, feature parameters around storage and simultaneous recordings can be changed and updated easily by utilizing the software control plane. New rate codes control what each subscription tier will look like. A RS-DVR platform can support multiple versions of a DVR product with different storage allocations and tuners. Upgrading storage from 160GB to 500GB or even to 1T can be done by control software only. The same is true for the number of virtual tuners. Since the tuners are not a limitation of the hardware, changing the number is just a variable that is used in the conflict resolution software. Similar to the storage flexibility, simultaneous recordings can be set to four, six, ten, or any number just by utilizing software in the headend.

For example, the initial Cablevision RS-DVR product launched with four tuners and 160GB of storage for the home. As always, any product decision about these values must be carefully vetted to completely understand the impact on the scaling of the complete system. This will be covered in detail later in this paper under the Capacity Planning section.

Whole Home Capability

Moving the recordings to the cloud also makes it possible to provide a whole home solution without adding additional components in the home. When the content is in the cloud, consumers have the potential to access the content from any device in the home. This is not limited to Set-Top boxes as it is now technically possible to deliver this content to any supported device in the home. One could design the system to allow a customer to start watching on her big screen and continue on her iPad.

Any devices in the home could become a DVR by downloading new application software and being authorized for the service. The devices continue to function on the network and do not require any other updates.

CLIENT CODE ARCHITECTURE

In order to enable cloud based recording and playback, client code is needed to utilize the remote systems. There are two areas within the user interface that need to be enhanced to support this architecture, the channel guide and the DVR UI that lists schedules and recordings.

In the channel guide, all record requests are intercepted and directed to the cloud-based RS-DVR backend. For example, if a customer is on the guide and navigates to a program, then selects to record a show from the guide, the information about the program as well as any control parameters that the user might set is sent to the backend software control plane.

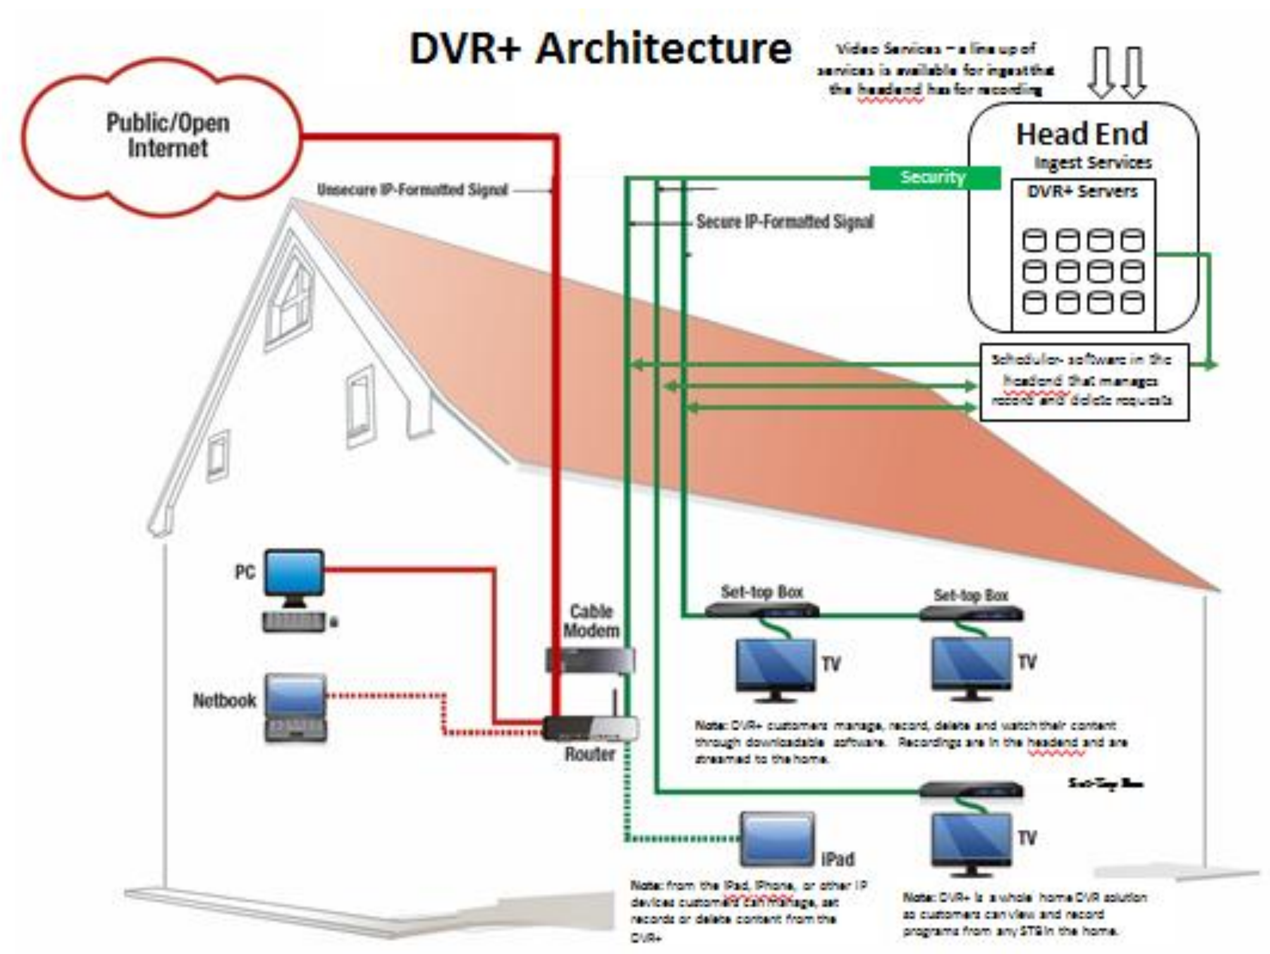

The DVR button on the remote control is programmed to navigate the users' listings of recordings. The RS-DVR guide calls back to the network to retrieve the list of recorded and scheduled programs. Since the DVR guide is in the headend it can be rendered by any STB associated with the account as well as by other IP devices with an Internet connection and proper credentials. Similarly any changes or additions to the requested recordings, new recordings, or deletes are reflected across all devices when rendering the recorded and scheduled lists.

An API layer exists which exposes interfaces to the RS-DVR client application(s). These APIs include functions such as:

- getScheduleList()
- getRecordedList()
- scheduleRecording()
- deleteRecording()

(These are just a few)

These APIs are available and used by all supported devices in the RS-DVR system.

SOFTWARE CONTROL PLANE

Scheduling

The component in the headend that brokers recording requests on behalf of the subscriber is called the Scheduler. The Scheduler is responsible for delivering 'just-in-time' requests to the physical recording system where the subscriber recording occurs.

The Scheduler logically sits between an application on the client device and the systems that control ingest and recording. (Figure 2). Management commands come into the system asynchronously throughout the day. These commands are managed by the Scheduler which, in turn, inform the ingestion service when to start a recording. A well-defined interface between Scheduling and Recording server must exist for this interchange. On successful execution of the recording, the system will commit the recording metadata for display to the subscriber when the RS-DVR menu is invoked.

Series Expansion

This Scheduling component also has an integration point in the headend with guide data. The operation of making recordings is no longer dependent on the STB or even the STB being connected to the network. The Scheduler will use guide data in the headend to make future recordings for series content based on subscribers' requests.

The process of ingesting daily guide data and scheduling recordings based on series settings of the subscriber is a function of the Scheduler. This process is referred to as "series expansion". The function of making recordings is also logically abstracted from the availability of the content RF signals in the home. Details of the ingest services will be discussed in the next section; however, for the purposes of scheduling it is important to point out that content signals are captured at the headend not at the home.

Entitlements

When the recording function moves out of the box and into the headend so does the need to do entitlement checking. Once again the architecture logically as well as physically de-couples the entitlement function from the in-home device. In the in-home DVR, if a subscriber sets a future recording for content they are not authorized for, the STB will tune to the channel and record black (or some barker) since the tuner is not able to decode the signal on that channel. In the cloud all channels are available at the ingestion service, so the system needs to check the subscribed status of the requested content for each subscriber. This software component is part of a standard service delivery platform that is also used for all server-side authorization checks.

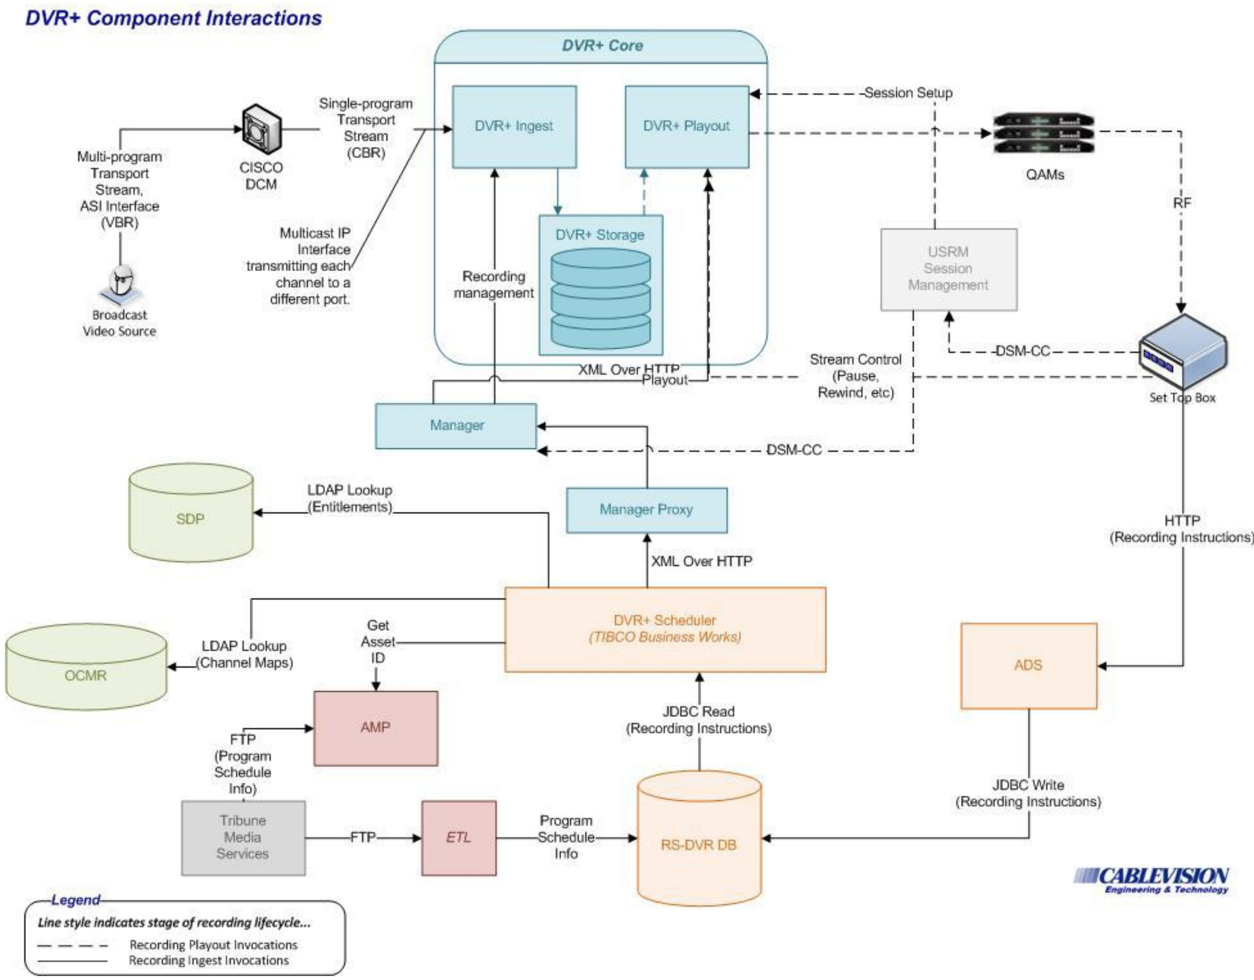

Figure 2 – Reference

INGEST STORAGE AND STREAMING

A main component of RS-DVR is the system that ingests stores and streams the content. (Figure 2. RS-DVR core)

Ingest

Each request received from the Scheduler represents an individual recording command from a subscriber. The ingest process of the RS-DVR system is responsible for ingesting the data/video associated with the program and storing it on behalf of the subscriber.

To enable this, the complete channel lineup must be available to the ingestion process with each channel having a unique multicast IP address assigned to it. The video quality and proper format is critical within this delivery portion of the system.

This component is the network equivalent of the application, in disk based DVR boxes, that tunes the box to the appropriate channel to move the content from the decoder to the local disk. It must be performed in a highly scalable manner as it occurs on behalf of every recording currently scheduled by the system. The data is read from the multicast IP and written to the storage location allocated to the subscriber(s).

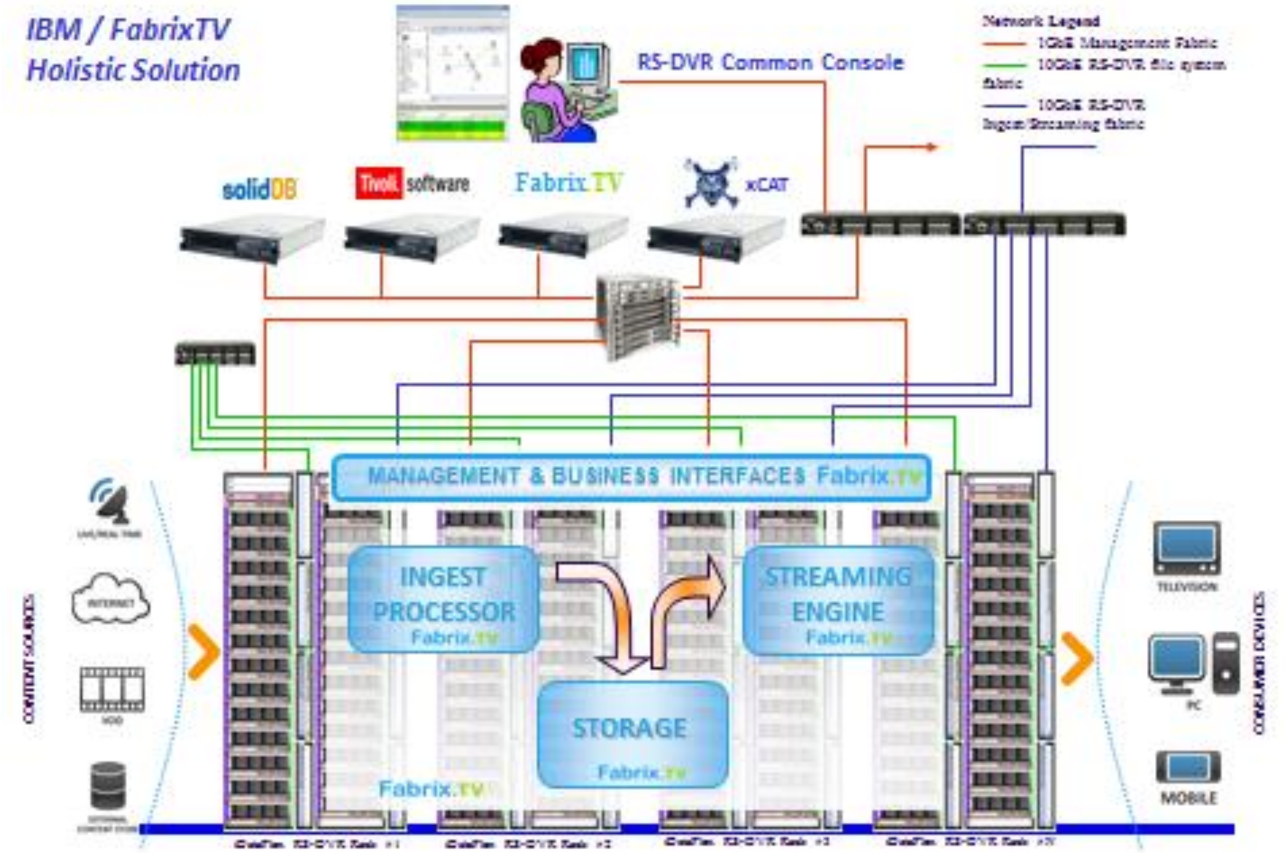

Figure 3 – Operational System Storage

This system demands individual program recordings for each subscriber and requires the right amount of storage that performs well in a write intensive environment. This posed challenges around finding a solution that provides high performance and does not require proprietary hardware. Commodity off the shelf, (COTS) servers and disk storage with the capability to access the disk drives to their maximum level is the solution that works best for this type of system.

Combining the use of a grid-based software file access system provides the best performance to extract the most out of the disk drives.

CAPACITY PLANNING

Building the components described are only a subset of the areas to think about when building the cloud version of DVR. In order to service all requests for recording and playback during peak simultaneous usage of the system it must have enough total throughput capacity in the core. The following is a list of parameters to consider when defining the system scale.

Variables for Storage:

- Total number of subscribers
- Storage allocation per subscriber
- Total available system storage
- Oversubscription percentage

Variables for Ingest:

- Total number of subscribers
- Expected concurrency

System Consideration Parameters

- Disk capacity
- Retro Fit ("Y"/"N")
- Max streaming concurrency
- Max ingest concurrency
- Storage Per Sub Offered
- HD/SD Ratio
- Simultaneous Peak Streams AND Peak Recordings
- HD Rate
- SD Rate

- Homes per Rack (POD)
- Usable Home Capacity Product Assumption
- Total Rack Throughput Capacity

Assumptions about the usage and product features for the service are used to calculate the system capacity needed to support the customer base. As an example assume the following parameters:

- Subscribers are using 100GB of storage on average
- Concurrency rate at peak usage time 10Mb/s / subscriber

To calculate the number of subscribers that can be put on a reference rack that has 28 servers, 12 drives/server, with drive size at 3T one would use the following:

Calculations of System Requirements DVR+ Core in the Reference Architecture
To go from a single POD to sub count: **Drive Size * No of Drives/server * No of Servers/POD * %age (RAID reserve, cushion, et)** **3TB * 12 * 28 * .85 = 856.8 TB usable storage/POD**
To determine subs/POD based on storage estimates**:** **POD Storage size/Storage Size/Sub (estimated)** **856.6/100GB = 8560 subs/POD**
To calculate total subs to no of PODs based on storage, it's just: **No of Subs*estimated size/POD Storage Size,** **300000 subs * 100GB / 856TB = ~35 PODs**
For ingest capacity**:** **Server ingest capacity * no of servers/POD**
3GB/sec * 28 = 84GB/sec ingest capacity/POD
To determine subs/POD based on peak ingest concurrency: **84GB/sec / 10MB/sec (estimated concurrent) = 8400 subs/POD**
To go from total subs to no of PODs based on ingest**:** **No of Subs*estimated concurrent/POD ingest capacity:** **300000 * 10MB/sec / 84GB/sec = ~36 PODs.**

Other aspects of the system to consider in capacity planning are the following:

Facility Availability

Based on the current demands on storage a facility to house the storage system is required.

Power Requirements

Power in the Headend must be considered as part of the overall installation. In addition, the shift from in-home DVR to DVR Plus in the cloud shifts some of the cost of power from the subscriber to the operator.

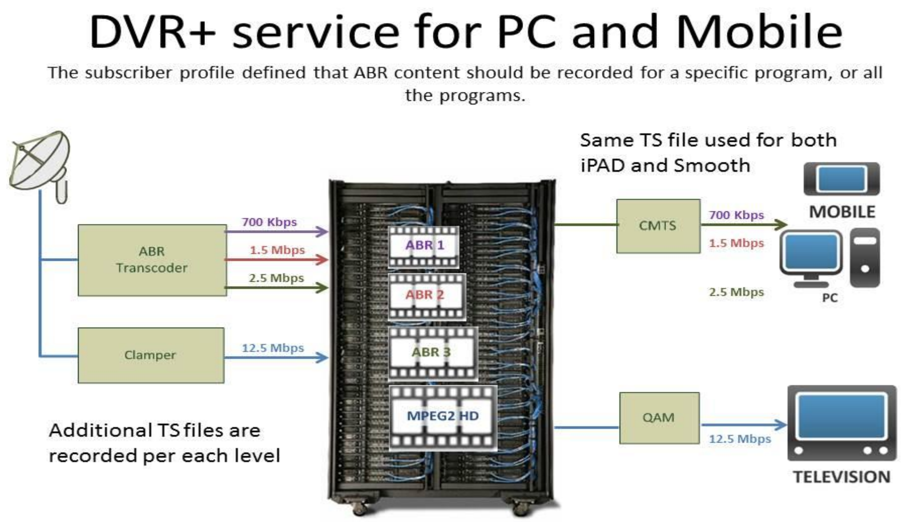

Figure 4 – cross platform distribution

Bandwidth Needs

Considerable network IP bandwidth is needed to deliver data both into and out of the RS-DVR Ingest, Storage and Streaming system. Playback of RS-DVR content to the subscriber uses the same network components as a standard VOD stream. QAM bandwidth is utilized similarly to VOD in that sessions are set up and allocated in order to stream the recorded content to the home. Utilizing a USRM facilitates sharing of the QAM resources between multiple uni-cast applications in this case VOD and RS-DVR. IP bandwidth inside the core network needs to be minimally built to the capacity of the availability of the QAMs. This calculation is based on number of QAMs * the number of service groups * the 38.8 mb/sec bandwidth available at each QAM.

TECHNOLOGY EXTENSIBILIY

The original purpose of moving DVR into the cloud was for recorded content. However, utilizing the RS-DVR cloud system is not limited to just DVR recorded content. There is value in using the cloud for other applications such as VOD, IP content distribution and advanced advertising. The technology component for enabling these types of content is the back office software.

The RS-DVR platform is enabled to support VOD content and library expansion that is not tied to traditional VOD vendor technology. Utilizing the scale of the storage system of the RS-DVR platform it becomes a natural extension of the VOD platform. There is storage available to support expansion of the library and streaming capacity to deliver the content. RS-DVR and VOD technically vary by the mechanism that is used to ingest the content into the system. RS-DVR uses customer-initiated commands to ingest and record content, while the content providers deliver VOD content by traditional pitcher/catcher exchange or over the network via FTP applications that deliver high-speed transfers into the systems. The new technology specific to the RS-DVR system helps in the turnaround time for VOD titles.

The system may also act as an origin server for IP delivered content. The same content sources can be transcoded and packaged within the cloud and available for subscriber requests. Additionally a CDN may be leveraged as the go between for the subscribers and the content.

A dynamic advertising solution is also possible in the platform through the use of SCTE-130 standards. Content ad markers for local insertion opportunities are detected on ingest. On play out, local ads may be inserted in the stream, based on responses from any desired Ad Decision System (ADS). For Cablevision's RS-DVR system, the system maintains linear parity. The ads are the same ads that were shown in the live broadcast within each of the zones. The ad copy is ingested into the RS-DVR core and is available to the streaming system for splicing into time-shifted unicast content as each individual playback calls back to the server for the recording. This advertising solution can be enabled to use any number of ADS systems as potential business emerges.

The original focus of engineering a dynamic advertising solution in the cloud was based on reducing cost and complexity of the overall system. The SCTE-130 solution was engineered to solve the challenge of recording content in multiple ad zones for subscribers in multiple regions. In order to preserve advertising in recorded content without DAI, it is necessary to take multiple versions of broadcast content and make them available to the ingest service of the RS-DVR core. To accomplish this it requires backhauling of the content from regional zones to the master headend. The cost and complexity of this is illustrated by taking the number of local ad zones and multiplying by the number of channels, both HD and SD, used in the commercial insertion system. As an example if there were 100 channels across 36 zones, 3600 video signals would have to be present at the ingest service location.

The DAI solution resolves the complexity and cost of backhauling local ad zones while maintaining parity with the Commercial insertion system. (figure 4.0)

An additional benefit of the cloud solution is the ability to continue to run regardless of the condition of the subscriber's home. For example, if subscribers home is affected due to a disruption services to their home due to a

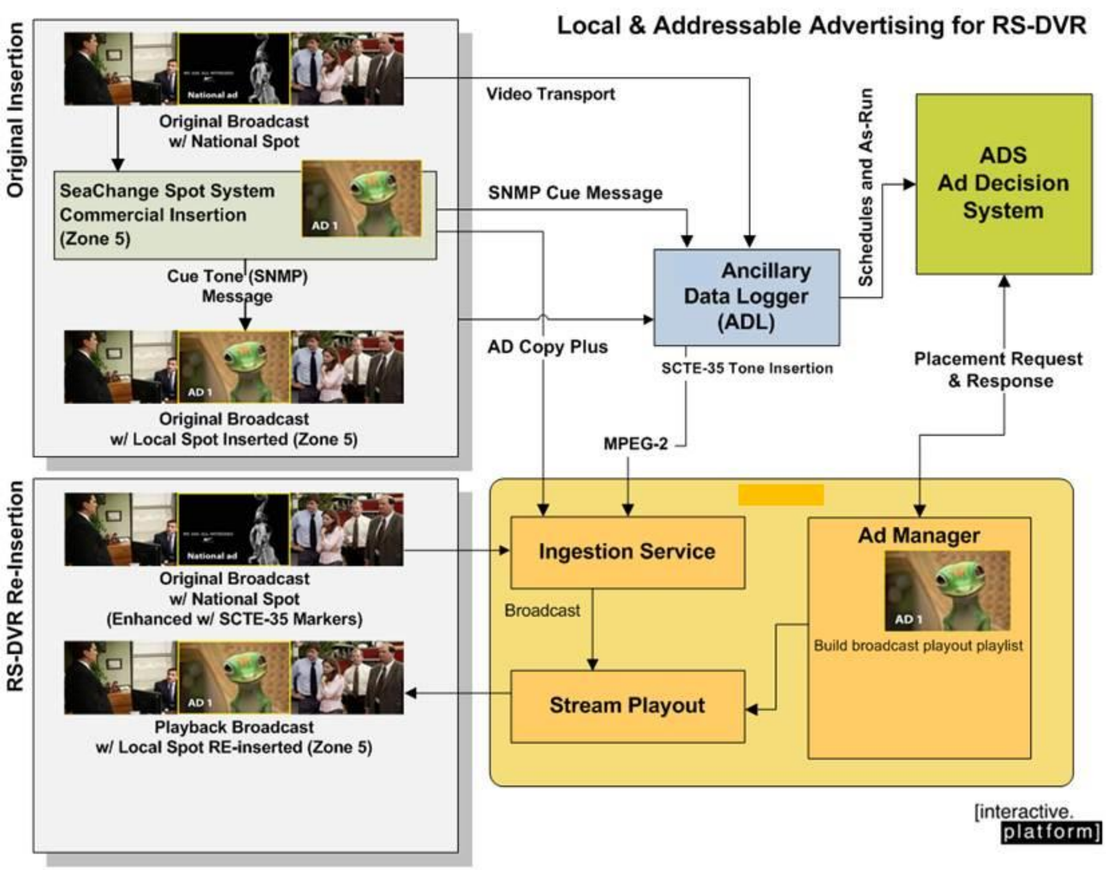

local power outage, their recordings will continue without disruption as long as the DVR Plus system remains online in the headend.

CONCLUSIONS

The idea of moving video services to the cloud is not a new concept. Internet companies are well aware of the fact that having video in the cloud has some strategic advantages. The most obvious one is mobility. Once the content is in a facility that can make it available to multiple devices the product opportunities are endless. Whole home and mobile DVR services are a clear differentiator for the cable companies.

Taking this system from the vision state to deployment is challenging. There are a number of educated guesses that must be made to model streaming bandwidth demands, storage requirements and ingest concurrency. A model is needed to set these values. This model however will only be as accurate as the assumptions that are used in it. After deployment monitoring the system has given us data to validate assumptions. As more subscribers are added to the system it is critical to view and analyze how users are using the product.

Operating this system for two years has given us a unique view into what is needed to support our customers. Since the system operates from the headend there is little for a customer service representative or field rep to do in the home as long as the device in the home, Set-Top box or IP device, is working properly and the home has connectivity. Monitoring for this service is focused on the headend and network components needed to deliver the service. We no longer have to ask the subscriber to return his/her DVR box to a walk in center when the in home disk fails.

Examples of Usage Data

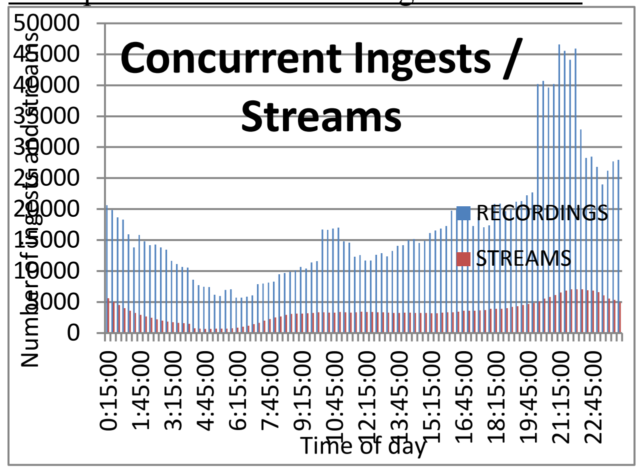

Figure 5 – daily streaming and ingest

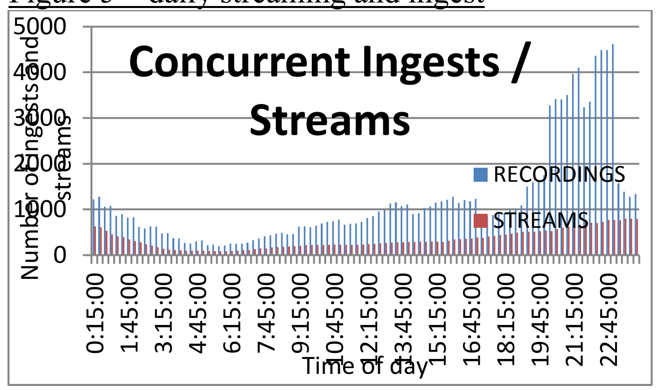

Figure 6 – daily streaming and ingest day 2

The data shown in Figure 5 and figure 6 show streaming and recording activities of a region during a 24 hour period. It shows how many recordings are being made on the system as well as how many streams are initiated. The second chart shows the same data just for a different day. It is clear from the charts that subscribers record more shows than they watch. This pattern is somewhat typical however the usage will vary based on a number of factors including the popularity of a specific show or series.

We must also consider how product features that are offered to the customer in the form of the User Interfaces, preferences, defaults and control affect the performance envelope of the overall system. For example when setting defaults it is helpful to have content delete set to after 14 days with the other options available but left to the subscriber to choose. Conversely we have found that even though we allow customers to record four shows at the same time very few actually do. Figure 7 shows the amount of content that is deleted over time

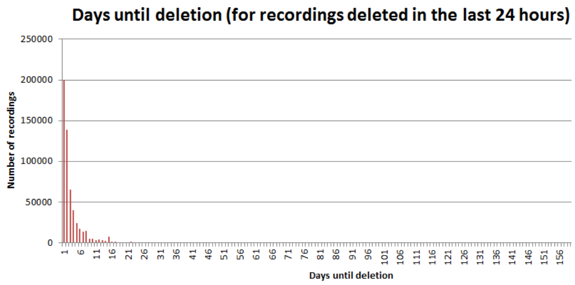

Figure 7 – content deleted over time.

You can see from the graph that the majority of content is deleted less than three days from the time it was recorded with the highest amount deleted in the first day. What this chart also indicates is the many customers keep the default setting for deleting content. It is difficult to see in this chart but at the fourteen day mark there is a noticeable pike in content deletions.

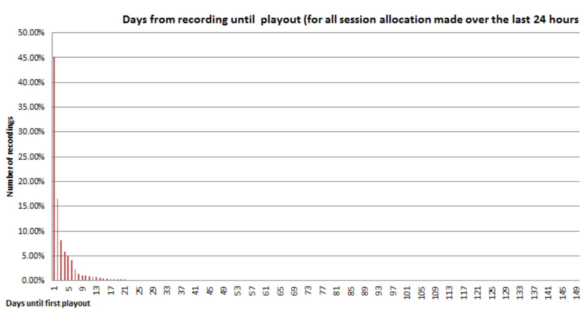

Figure 8 – days until content is watched after recording

This chart shows us that the majority of content is watch within two days from the time it was recorded. So the two graphs above tell us that average behavior for people using the RS-DVR system is that they record, watch and delete content in a relatively short period of time after the show airs.

Figure 9 is a historical view of our RS-DVR deployment. As you can see in the chart streaming and ingest rise at a consistent rate in time as we on board new subscribers to the system.

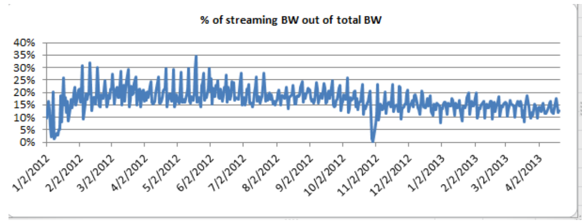

Figure 9 – System Streaming BW utilization

As figure 9 shows it is clear that over time and while the subscriber base is expanding the overall stream bandwidth utilization is low. This is because the system just by the nature of the product is a recording platform. Since the platform has streaming capacity available for other application it is the logical place to use for other content distribution products.

It is important to emphasize, that using a model that takes into account the parameters that most accurately describe the region, number of subscribers and product offering is critical to building out this platform. It is also important to understand the operational ecosystem that a RS-DVR system need to operate in.

This architecture enables the future video grid. It will support the vision of cross platform distribution of content in a manageable way and transform over time from a traditional MPEG2 ingest and storage platform, to a multi-format distribution system that will server up the right format for the client.

Three components of the future architecture are content acquisition infrastructure, a software control plane and a distribution network. Use of the described framework will prepare the industry to continue to deliver quality content and products over multiple networks to numerous devices. Most importantly client/customer owned equipment.

SMART ABR: THE FUTURE OF MANAGED IP VIDEO SERVICES

John Ulm, Ajay Luthra, Praveen Moorthy, Mark Schmidt, Haifeng Xu
Arris, Inc. (formerly part of Motorola Home)

Abstract

This paper documents the research and lab results for using Adaptive Bit Rate (ABR) protocols for a managed video service. It quantifies the issues with unmanaged ABR including unfairness, instability and inefficiencies. It then explores some potential solutions including using either CMTS QoS or Server based algorithms in the cloud.

This ABR research has led to the evolution of Smart ABR (SABR) allowing operators to provide a first rate video service with exceptional Quality of Experience while retaining the underlying benefits of Adaptive protocols. The paper highlights lab results showing the optimization and Video Quality achieved. With SABR, operators can significantly increase their IP Video capacity while gracefully handling congestion and providing an improved user experience.

INTRODUCTION

Adaptive Bit Rate (ABR) protocols have become the mainstay of multi-screen devices like tablets, smart phones, gaming devices and Smart TVs for accessing Over-The-Top (OTT) video content. Because of their explosive popularity, it is highly desirable for an operator to provide existing video services to these devices. However, consumers will expect the same Quality of Experience (QoE) to which they are accustomed with today's primary TV screen delivery.

The ABR protocols have been optimized to operate over an erratic internet connection. However, the ABR client based control with no insights into system behaviors has demonstrated many inappropriate behaviors. The ABR client's greedy behavior leads to significant unfairness, instability and inefficiencies. These traits are not suitable for an operator to offer a true managed ABR video service with the associated QoE.

Operators have a number of challenges in offering a "managed" ABR video service. The limited bandwidth makes it challenging to support the demand of a large number of concurrent users while maintaining good video quality for each user. In addition, existing implementations of ABR client controlled distribution mechanisms are not very efficient. They tend to under utilize the available bandwidth and provide uneven visual qualities to the clients. Therefore, understanding the issues around delivery of ABR over the DOCSIS network will be crucial for MSO's video service delivery, and for their ongoing profitability.

Research into ABR has led to the evolution of Smart ABR (SABR). By adding some cloud based intelligence back into the system, the operator can regain control to provide a first rate video service with exceptional Quality of Experience while retaining the key underlying benefits of Adaptive protocols. SABR is a server controlled system that can manage the bit rates and video quality that each client receives.

Intelligence in the SABR server maintains client state, available client download bandwidth or channel capacity, and a measure of "reasonable" client video quality. "Reasonable" could be dependent on client attributes such as client display size or the type of video content being watched. Based on that intelligence, the SABR server controls what bit rate each client gets. This can prevent oscillations in network bandwidth utilization and increase network utilization.

The paper discusses lab results showing the optimization and Video Quality achieved. The SABR system is compared in detail to traditional unmanaged ABR delivery as well as a system with enhanced CMTS QoS. With SABR, operators can significantly increase their IP Video capacity while gracefully handling congestion and providing an improved user experience.

ABR Overview

Adaptive Bit Rate (ABR) is a delivery method for streaming video over IP. Adaptive streaming uses HTTP as the transport for small video chunks of approximately 2-10 seconds each. This enables the content to easily traverse firewalls, and the system scales exceptionally well as it leverages traditional HTTP caching mechanisms in the CDN.

Adaptive streaming was developed for video distribution over the Internet. In order to deal with the unpredictable performance characteristics typical of this environment, ABR includes the ability to switch at chunk boundaries between different encodings of the same content. This is illustrated in Figure 1. Depending upon available bandwidth, an ABR client can choose the optimum encoding to maximize the user experience.

The server stores several chunk sizes for each segment in time. The ABR client predicts the available bandwidth and requests the best chunk size using the appropriate URI. Since the ABR client is controlling when the content is requested, this is seen as a client-pull mechanism, compared to traditional streaming where the server pushes the content. Using URIs to create the playlist enables very simple client devices using web browser-type interfaces. A more in-depth discussion of ABR video delivery can be found in [Ulm].

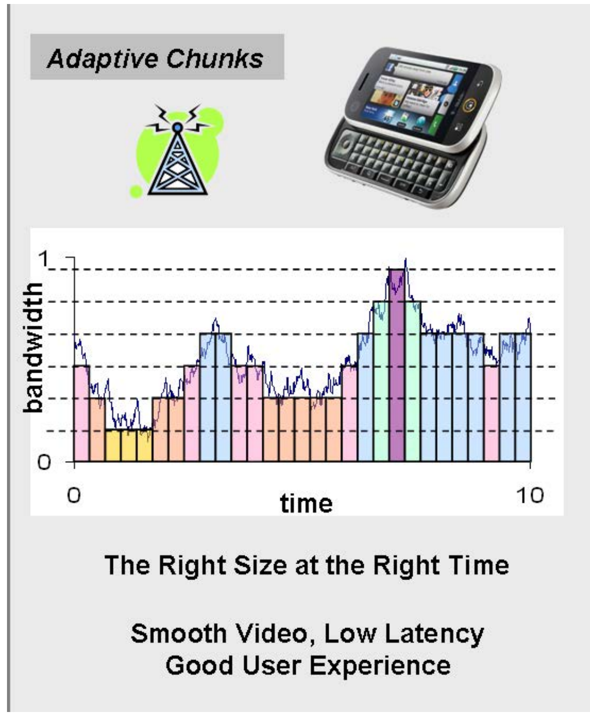

Figure 1 Adaptive Streaming Basics

Importance of ABR – 2^{nd} and 3^{rd} Screens

ABR based video streaming has become the de-facto standard for video delivery to IP devices such as PCs, tablets, smart-phones, gaming devices and Smart TVs. ABR clients are typically shipped with (or are available for download to) these devices as soon as they are released. Given the short lifetime of this class of device this is a key enabler, especially compared to the time required to deploy software to traditional cable STB devices. As mentioned previously, ABR delivery simply requires an HTTP connection with sufficient bandwidth so that it is available both on net and off net. With these advantages, essentially all video delivery to second and third screen devices uses this mechanism.

ABR vs. Current Managed Video Delivery

ABR video delivery has a number of very significant differences to both MPEG video delivery and streamed IP video delivered over Real-time Transport Protocol/User Datagram Protocol (RTP/UDP) as used in a Telco TV system. Foremost, ABR has been developed to operate autonomously over an unmanaged generic IP network.

- *The client device decides on bit rate (i.e. bandwidth) decisions based on <u>its</u> interpretation of network conditions.*

This is fundamentally different from the approaches used for existing MPEG or conventional streamed UDP video delivery, where devices under the direct control of the network operator make the important decisions relating to bandwidth. Thus, in MPEG delivery, the encoding, statistical multiplexing and streaming devices determine the bit rate for a given video stream. These devices are under control of the service provider. In contrast, the behavior of ABR clients is specified by the CPE developer which, in general, will be a third party outside the service provider's control.

An ABR client selects a file chunk with a bit rate that it believes to be most appropriate based on a number of factors including network congestion (as perceived by the client) and the depth of its playback buffer.

- *Thus the load presented to the network can fluctuate dramatically.*

Operators in a controlled network can guarantee that adding new user sessions do not impact existing users. Once resources are exhausted, any additional session requests will be denied, introducing a probability of blocking into the system.

ABR clients join and leave the network as users start and stop applications. From a network perspective, there is no concept of a session with reserved resources or admission control. Again this is the antithesis of MPEG or UDP video in which the control plane operates to request and reserve network resources and determines when to admit users. In a pure ABR model with network congestion, each new session will reduce the bandwidth available to all existing sessions rather than be denied.

- *Thus, all users may see a variation in video quality as other ABR clients start or change bit rates.*

With MPEG or UDP streaming video delivery, congestion control is not relevant as the control plane provides admission control to ensure it does not occur. When ABR is used for video delivery, congestion control is a potential issue. The situation is complex in that three levels of congestion control mechanisms are involved operating at different layers in the protocol stack. At the media access control (MAC) level, the CMTS is responsible for scheduling downstream DOCSIS traffic. Operating at the transport level is standard Transmission Control Protocol (TCP) flow control based on window sizes and ACKs. Finally, at the application level the client can select the video bit rate to request. The latter two levels of control (TCP and application) are the responsibility of the ABR clients and as such are outside the control of the network operator. Interaction between these flow control mechanisms is not well documented and may have unforeseen impacts.

In summary, ABR clients base their decisions on what to request based on their local knowledge and observed conditions rather than on an overall view of the network conditions. This is in contrast to MPEG or UDP streaming where the network operator provisions the video bit rates based on knowledge of the end-to-end network and expected loads.

Adaptive Bit Rate streaming is deployed today in a number of implementations, including MPEG Dynamic Adaptive Streaming over HTTP (DASH), Apple HTTP Live Streaming (HLS) and Microsoft HTTP Smooth Streaming (HSS). Due to the popularity of HLS and the abundance of iPad and iPod devices available for our testing, HLS was used in our tests. The techniques and solutions described in the paper may be applied to other ABR formats that rely on HTTP delivery of segmented content.

POTENTIAL MANAGED ABR SOLUTIONS

ABR Standards Investigations

The MPEG DASH Ad Hoc Group has been investigating the use of quality-driven streaming in a DASH environment. In this approach the client is provided with additional video quality information that the client can use during its segment selection and rate adaptation process. This info is generated during content preparation, converted to an estimated video quality measure and carried with each Media Segment.

While early experiments have shown improved QoE, MPEG DASH does not specify a normative client implementation or behavior, so the full benefit of client-directed, quality-driven streaming is dependent on well behaved clients using similar algorithms for segment selection/rate adaptation.

While this approach may improve QoE for a given client, it is not clear that this will solve the collective issues of unfairness and instability across a group of clients as seen in unmanaged ABR. Also, operators must offer a video service across a wide range of clients that must include others besides DASH. For these reasons, DASH was not considered as a managed ABR solution.

CMTS as Control Point

For users on an HFC network, IP traffic will always flow through the same CMTS port to reach a user at home. As the shared CMTS to CM link is normally the “narrow pipe” in the video distribution network, this is where congestion would be expected. Therefore the CMTS can potentially provide a useful control point to manage ABR traffic.

The DOCSIS standard provides very complete Quality of Service (QoS) functionality which may be useful for managing ABR traffic. If a packet matches an installed classifier it will be mapped to a specific Service Flow and then forwarded based on the parameters associated with that Service Flow. Classification is based on matching fields in the packet header such as Source &/or Destination IP address, port & type and Differentiated Services Code Point (DSCP) fields.

Therefore it is possible to recognize a managed ABR video packet stream from a well known source address (e.g. video server) or IP subnet. The CMTS could then provide preferential QoS treatment for the operator’s managed video flows. One of the goals of our research was to verify the impact of using CMTS QoS for managed ABR video service.

The DOCSIS infrastructure has a mechanism to dynamically setup and control Service Flows based on the PacketCable™ Multimedia specification [PCMM]. This provides a potential mechanism to implement resource reservation at the session level. It requires a session establishment and teardown mechanism. However, this may be problematic with the distributed nature of adaptive streaming protocols and may have scaling issues. The control plane topic is outside the scope of this paper.

Cloud based Server Solutions

In conventional ABR video distribution, the ABR client determines the bit rate of the next file to download from the options in the playlist and retrieves this directly from the Content Delivery Network (CDN). By adding intelligence into the cloud, this decision could potentially be overridden from the network in a number of ways. This is referred to as Smart ABR or SABR.

The playlist file provides the bit rate options specified by the service provider. Normally this selection would be statically provisioned and implemented by the encoding and packaging processes as the video asset was processed. It is conceivable that the playlist may be manipulated and the network can regain control of what bit rates are available to the ABR client.

During peak utilization, existing managed video delivery uses admission control to block new users from accessing the system. With SABR, an operator could gracefully handle congestion during peak times with no blocking of users, but rather a slight degradation of video quality. This reduction in quality during peak times is analogous to statistical multiplexing in legacy MPEG video. During peak times, the statmux reduces bit rates across the various video streams to fit within its channel. The SABR system has an advantage in that it will be multiplexing over a larger channel using DOCSIS bonding.

In a SABR system, all the clients are controlled from the server side. The system level intelligence in the server understands the state for every client; the available bandwidth for each client; and a "reasonable" visual quality of the video for a given size of display and attributes of video etc. Based on that intelligence the server controls what bit rate each client gets. This avoids the oscillations and increases network utilization.

ABR TEST METHODOLOGY

Our goals were to research ABR behaviors in a working environment. First item was to replicate and quantify existing unmanaged ABR characteristics. Then test these same conditions for a CMTS QoS based solution and a cloud based SABR Server solution.

To analyze our lab results, it is important to understand some of the fundamental operation of the ABR client.

ABR Client Characterization

The ABR client plays a critical role in the operation of adaptive protocols. For an operator trying to provide a differentiated quality of experience, it is important to understand how different ABR clients behave under various circumstances.

Previous work [Cloonan] discussed results from a simulator. Our goal was to capture live client interaction. Operation during steady state was relatively stable.

HLS Client Model

Because of the abundance of iPad and iPod devices available for our testing, HLS devices are used in these tests. A simplified HLS client diagram is illustrated in Fig 2.

A stored HLS program such as Video On Demand (VOD) assets have a manifest that lists all of the program's available media chunks or segments and the player downloads chunks starting from the earliest. When the client plays a stored program, it first reads a manifest file (playlist) from HLS server with specified URI, parses the content of the file, and starts to request HLS chunks sequentially starting with the lowest sequence number.

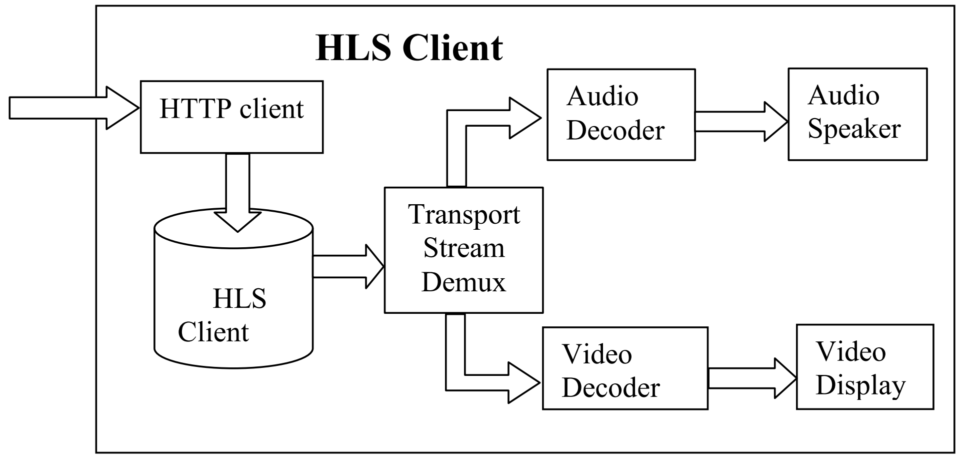

Figure 2 Diagram of HLS Client

A live Linear HLS program's manifest changes as new content are created; a sliding-window of chunks is given to the player and the player may or may not download the earliest chunk in this manifest.

The video play back does not start until the buffer hits a certain threshold. Therefore, if a client is not able to fill its buffer fast enough to a certain designed level, it will take longer to start the play back. Once the client fills up its buffer, it moves into the Playback stage.

In the Playback stage, the client fetches one HLS chunk during each chunk period. In other words, steady state is achieved and the overall download chunk speed matches its real-time play speed.

The two phases are illustrated in Figure 3. As can be seen, a client puts higher stress on the network bandwidth during the Buffering stage than that in the Playback stage as the clients try to buffer multiple segments as fast as they can. This introduces the following inefficiencies:

- *The HLS client buffer is necessary to deal with network jitter and varying bandwidth. Therefore, it requires more overall bandwidth during Buffering stage to provide this cushion.*
- *The HLS client relies on a combination of TCP/IP mechanisms at the low layer and adjusting video bit rates at application layer to deal with network variations. To provide the visual quality that does not vary too fast, a HLS client will not utilize full network capacity.*

As a result of these, we have observed clients in our labs that may leave up to 50% of the available network bandwidth unutilized Note: this number will vary depending on specific HLS client implementation.

During this startup period, the clients are also calculating the available bandwidth and may decide to switch bit rate. This action may cause some segments to be re-fetched with the new resolution. Overall, the differences between clients seemed fairly subtle for startup.

Once HLS clients retrieving a VOD playlist content reaches steady state Playback stage, it may have 50-60 seconds worth of content in its buffer. Live content tends to have a limited playlist available to the client, preventing large buffer build up.

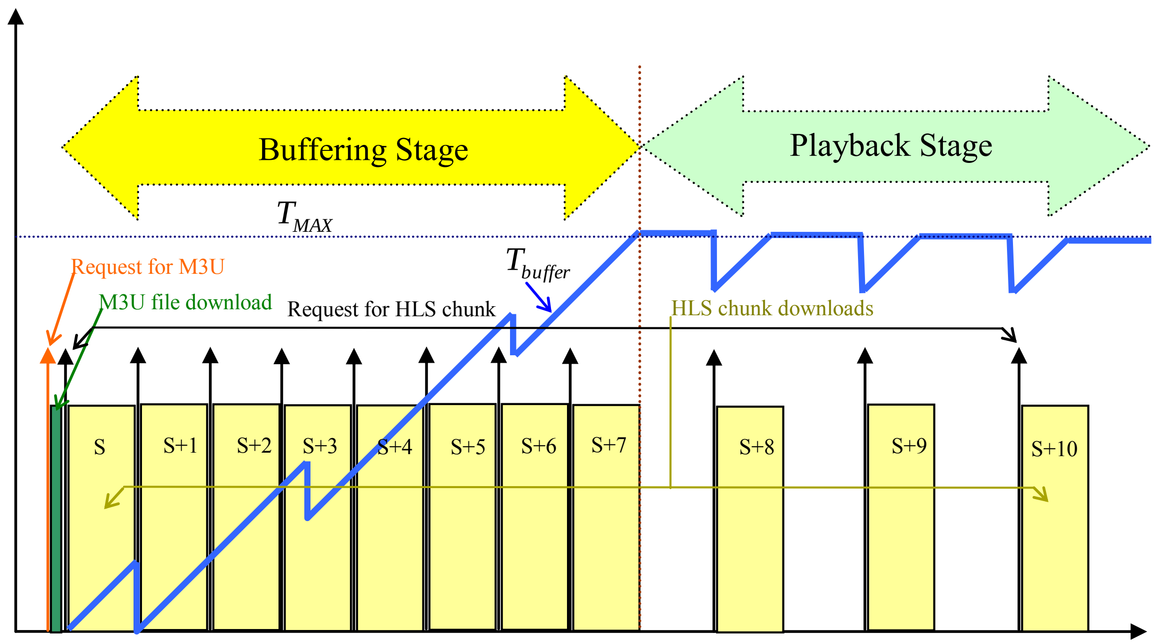

Figure 3 Two phase client stages

During playback, if an ABR client detects a significant enough decrease in available network bandwidth, it decides to switch to a lower available bit rate. This causes the client to move back to the Buffering stage and the client often downloads multiple chunks of the lower rate stream that correspond to chunks it has already downloaded from the higher rate stream. This could allow the client to seamlessly stitch video and/or audio in its decoder buffers. This download overlap may be as little as 2-4 seconds or as much as 12 seconds but still represents increased network load.

Conversely, when a client detects that its available download bandwidth has increased sufficiently, it switches to a higher bit rate and re-loads chunks corresponding to already downloaded media of the lower bit rate. The client is once again in the Buffering stage. However, these downloads can be substantial and clients have been observed to re-download enough segments to refill part or even the entire 50-60 sec buffer described above. Also, the client may quickly ratchet through multiple video rates as it tries to determine the optimum rate. These behaviors are significant contributors to the bandwidth oscillations among ABR clients.

Multiple Client Interactions

Based on the above characterization, operators must be aware of some potential problems. As was discussed, there is a burst of additional traffic during startup and when switching bit rates. For managed video delivery, the overall system must be capable of handling this additional traffic burst.

Actively managing ABR video traffic may be challenging given that every ABR client may be operating its own disjoint algorithm. This is also compounded since client behavior may change with the download of an updated revision. Bandwidth stability may become a concern if multiple clients become synchronized.

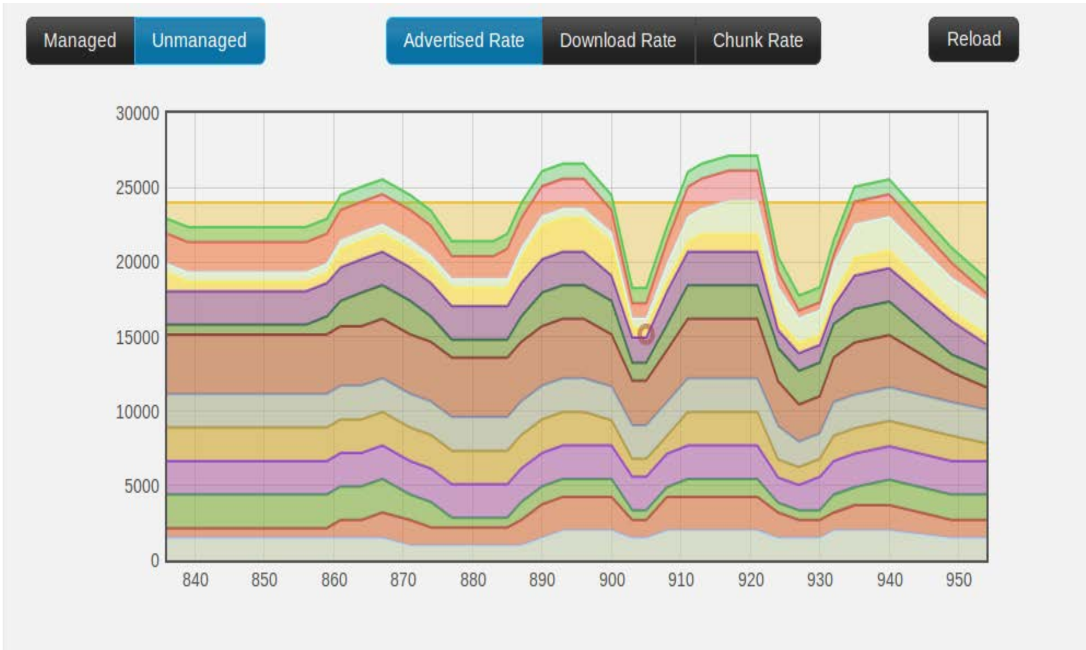

Figure 4 Multiple Client ABR Oscillation Example

For example, the network becomes congested causing a group of clients to lower bit rates. If these clients then sense that bandwidth is available (i.e. it is released due to downshifting by other clients), there may be a surge in traffic that causes congestion, and the cycle repeats. This oscillatory behavior is demonstrated in Figure 4 which shows the ABR bit rates versus time of thirteen HLS clients in a stacked bandwidth plot. The frequent stream bit rate changes results in changing video quality at the clients as well as inefficient network utilization. These issues of stability and fairness have been discussed in other papers such as [White] and [Adams].

ABR TEST SETUP

Our lab test setup needed to allow tests to replicate and quantify the issues with unmanaged ABR delivery and then show the impact of managed ABR solutions. A block diagram of the ABR test setup is shown in Figure 5. The test setup was explicitly designed to show a number of variations. More details on the lab test setup can be found in the appendix.

In these tests, there were only HLS video streams present and no other background Web data or VoIP traffic in the system.

Network Configuration

The overall network configuration consisted of an HLS Server, a CMTS, a half dozen cable modems with WiFi and then ten iPads/iPods. Since a primary goal was to observe behavior under network congestion, the CMTS was configured to a single downstream channel operating at 64-QAM. This provides approximately 25Mbps of user bandwidth. The Upstream channel was configured to operate at 10Mbps, sufficiently large that it shouldn't impact test results by delaying TCP ACKs.

A HTTP server hosted multiple ABR VOD media streams and was connected behind the CMTS along with a PC server running DHCP and TFTP servers used for cable modem configuration.

The six cable modems were dual-band 802.11n WiFi capable. Four cable modems used the less congested 5 GHz WiFi bands while iPods were associated with cable modems on the 2 GHz WiFi band.

ABR Video Clients

The test setup consisted of 7 iPads and 3 iPods clients. The number of clients per cable modem was intentionally varied from 1 to 3 to measure the impact of having multiple streams per home.

Apple iPads of various generations and HLS client software, iOS, was selected to observe if there were any impacts from different generations of iOS versions.

The clients were started in sequential order as listed in Table 1 in the Appendix. No attempt was made to allow each client to complete its buffering phase so it was expected there would be high offered load as the client's began their downloads one after the other. The tests lasted at least 20 minutes during which time logs were maintained of the advertised and actual bit rates of the HLS media segments (chunks) retrieved by the clients, the associated times of those chunk requests, the downloaded chunk sequence number, and the video quality.

Video Content

The HLS video streams varied in content complexity and resolution. The tests include both VGA (480p30) and HD (720p30) resolutions. Video clips were chosen such that a metric (i.e. PSNR data) was available for measuring video quality. A detailed description of how PSNR (Peak Signal to Noise Ratio) is calculated is in the Appendix.

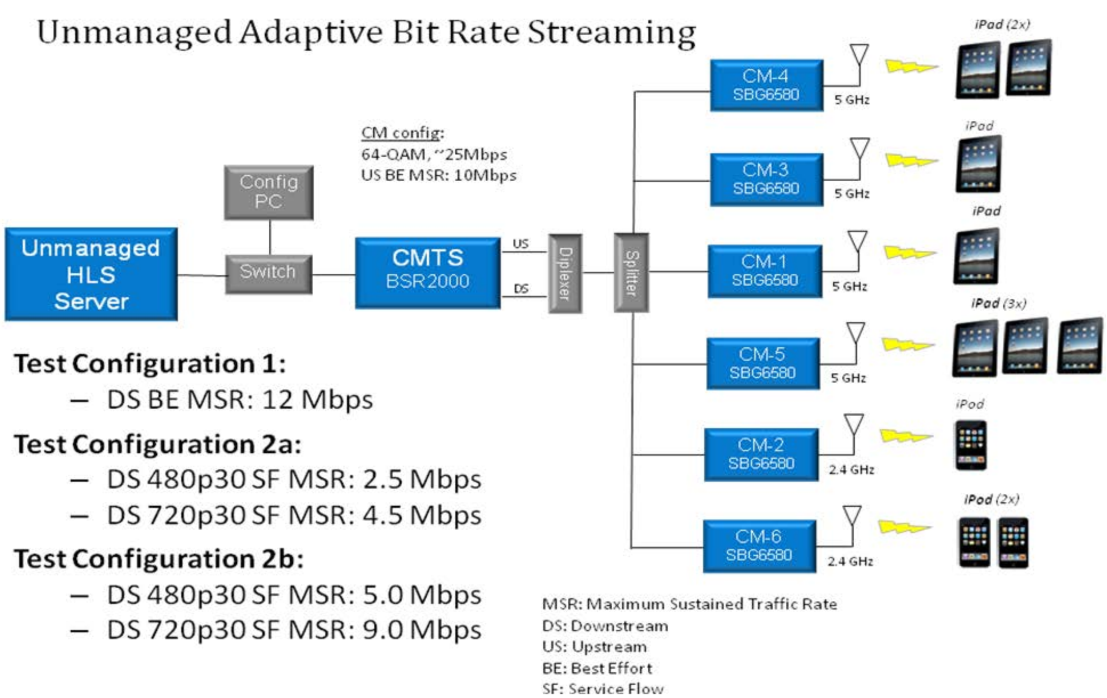

Figure 5 Adaptive Bit Rate Streaming Test Setup

The 480p30 VGA content supported four bit rates at 650kbps, 1.2Mbps, 1.7Mbps and 2.2 Mbps. The 720p30 HD content supported four bit rates at 1.0, 2.0, 3.0 and 4.0 Mbps.

CMTS QoS

For the unmanaged ABR tests and the SABR Server tests, the cable modems were configured for simple best effort traffic with a burst rate of 12Mbps. This rate was chosen to be less than the channel capacity while sufficiently large to support multiple video streams per cable modem.

For the CMTS QoS tests, each video client had its own Service Flow statically configured. One test was run with the max burst rate slightly higher (i.e. ~12%) than the HLS client's max video rate, then the remaining CMTS QoS tests were run with the max burst rate set to twice the max video rate (i.e. 100% extra Burst capacity).

UNMANAGED ABR TEST RESULTS – DEALING WITH A DRUNKEN TEENAGER

While analyzing the unmanaged ABR test results in the lab, an observation was made that the ABR clients were like dealing with a drunken teenager. Both are very erratic, self centered, greedy, unpredictable and hard to manage. This causes a slew of problems that are detailed below. The most intriguing observations occurred during startup and when video bit rates were forced to change.

The results of one of our unmanaged ABR tests are shown in Figures 6a, 6b & 6c. Figure 6a shows the aggregate bit rate from all 10 video clients for both the bit rate selected and the actual bandwidth consumed. The difference between these two lines is the result of clients being in the Buffering stage and requesting additional chunks to fill its buffers. Figure 6b shows a group of four charts depicting the instantaneous bit rate that each of the 10 HLS clients are requesting. The first, third and fourth charts show clients grouped by the cable modem that they share. The second chart is a collection of the three video clients with their own cable modem. Fig 6c is a collection of four charts showing the estimated video buffer depth for each of the 10 HLS clients. The four charts are organized in the same manner as in Fig 6b. [Note – this same layout is then used for the CMTS QoS results (Fig 7a-c, 8a-c) and the SABR results (Fig 9a-c).

Network Utilization

Our first observation is the network utilization from Fig 6a. The aggregate bandwidth consumed is very close to the 25Mbps capacity of this DOCSIS downstream. The average bandwidth consumed after that initial startup is ~23.6Mbps. The second line in Fig 6a shows the aggregate of all current bit rates selected. This averages about 19.2Mbps after the initial startup. This means that the unmanaged ABR scenario utilizes ~77% of the total channel capacity. The ABR clients have improved over the years. Early generation software was only seeing ~50% network utilization.

Instability

The next observation is on the stability of the system. In looking at the requested bit rates in Fig 6b, half of the video clients appear fairly stable after a couple minute startup period while the other half are oscillating for the entire test. These 5 HLS clients averaged between 1 and 1.5 bit rate changes per minute after the startup period.

Looking at Fig 6c gives an insight into the impact on this instability. The 5 volatile HLS clients had trouble maintaining their video buffer depth. This volatility would have been disastrous for live linear content.

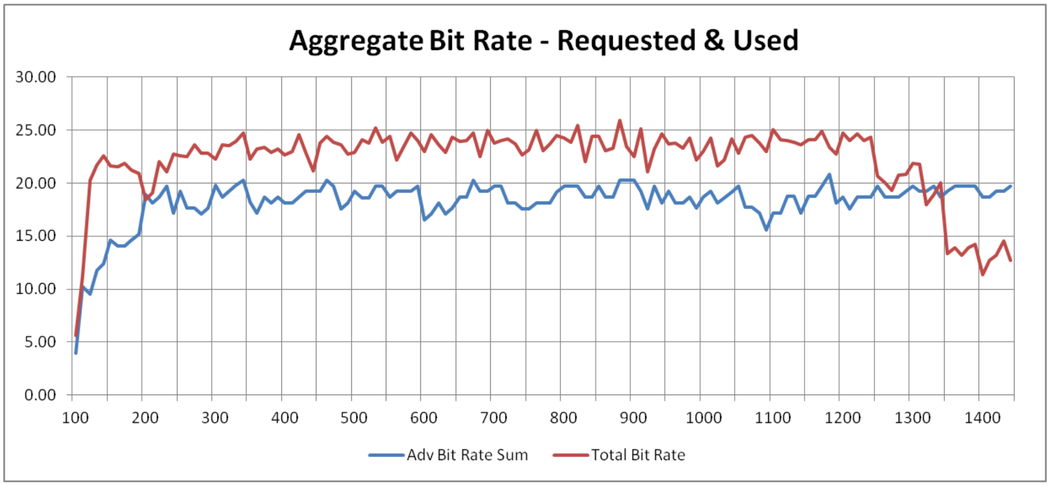

Figure 6a Unmanaged ABR – Aggregate Rates

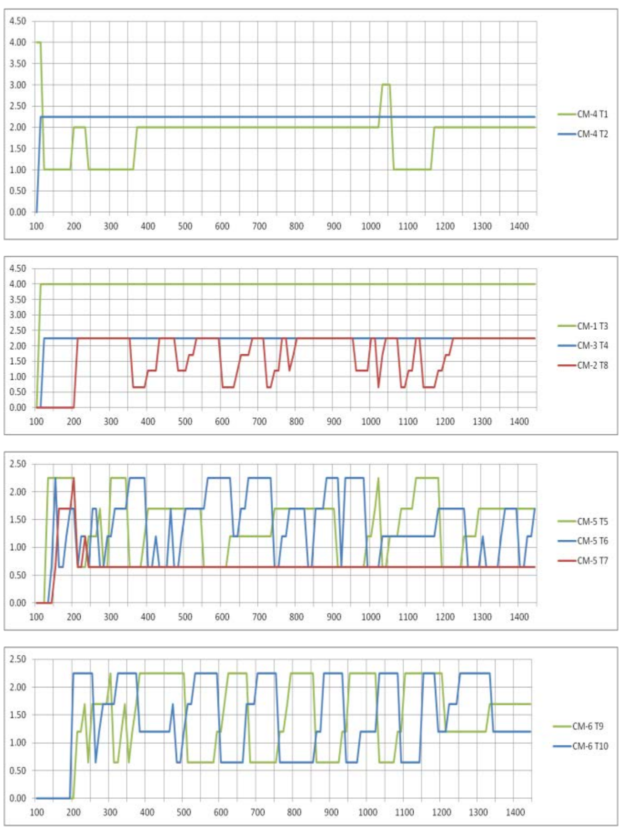

Figure 6b Unmanaged ABR – Bit Rates

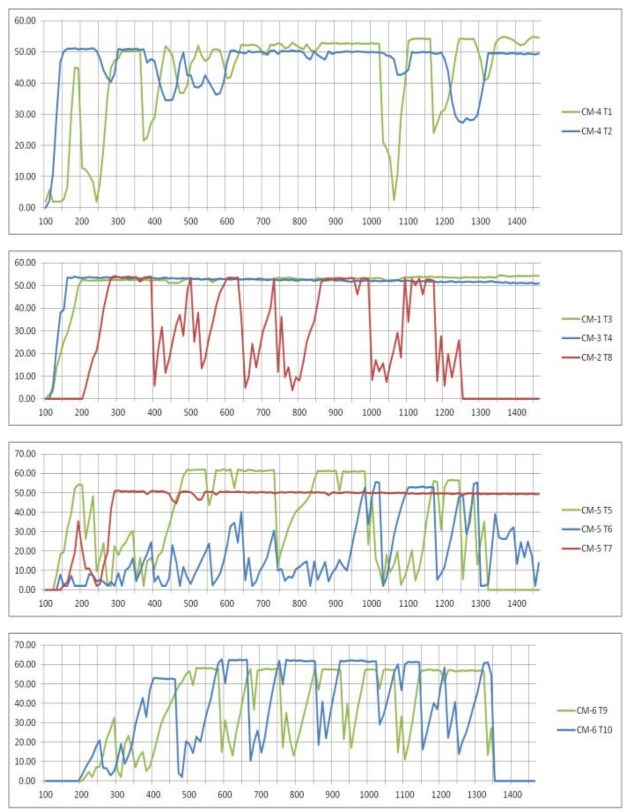

Fig 6c Unmanaged ABR – Buffer Depth

Unfairness

A key attribute for managed video service is maintaining fairness across clients. In running multiple unmanaged ABR tests, we observed multiple types of unfairness.

While both HD clients appear stable [see Fig 6b, CM-4 T1, CM-1 T3], one settles in at its maximum bit rate of 4Mbps while the other sits at 2Mbps most of the time. Similarly, two stable VGA clients settled at its max rate of 2.2Mbps (i.e. higher than one of the HD clients!!) and a third stable VGA settled at its min rate of 650Kbps.

To add insult to injury, one of the VGA clients at 2.2Mbps was a low complexity News clip (i.e. talking heads) while the HD clip sitting at 2Mbps was a high action and complexity Football sequence.

Other forms of unfairness may be introduced when network congestion causes video bit rate changes. Some clients may decide to change while others remain at current bit rates, resulting in disparity between clients.

It was also observed that in general the modems with multiple HLS clients struggled more than modems with a single client. The CMTS is trying to distribute bandwidth fairly among cable modems, so a home with three HLS clients find themselves competing for that home's bandwidth while another client may have the home's bandwidth all to itself.

We also observed that some HLS clients with older software struggle more compared to the clients with newer software. While overall utilization has improved over the years, this has apparently been done while making newer clients more aggressive and taking advantage of clients with older software.

In summary, we observed fairness issues with each of the following:

- *Client vs. client*
    - *Screen size (HD vs. VGA)*
    - *Video Complexity (Low vs. High)*
- *HLS client versions*
    - *newer more aggressive*
- *Multiple clients per home*
    - *WiFi impacts as well*

Video Quality

Some qualitative observations were made on perceived video quality while running the tests. The clients with stability issues would experience frequent changes in bit rates that resulted in reduced viewer experience. The changes in resolution were obvious. Some clients were oscillating from the max rate to their min rate. It appears that it would be better QoE if the rate stayed constant at a lower value then continually adjusting between different rates.

Even for the stable clients, the poorer resolution was noted for the HD client stuck at the lower rate and the VGA client stuck at its lowest rate.

Perhaps the worst impact on QoE was the buffer under-runs that cause the screens to freeze. Looking at Fig 6c shows how a number of clients struggled to keep an adequate amount of video in its video buffer. Two clients had aggregate buffer under-runs of 5-10 seconds and one client had over a minute of time with buffer under-runs.

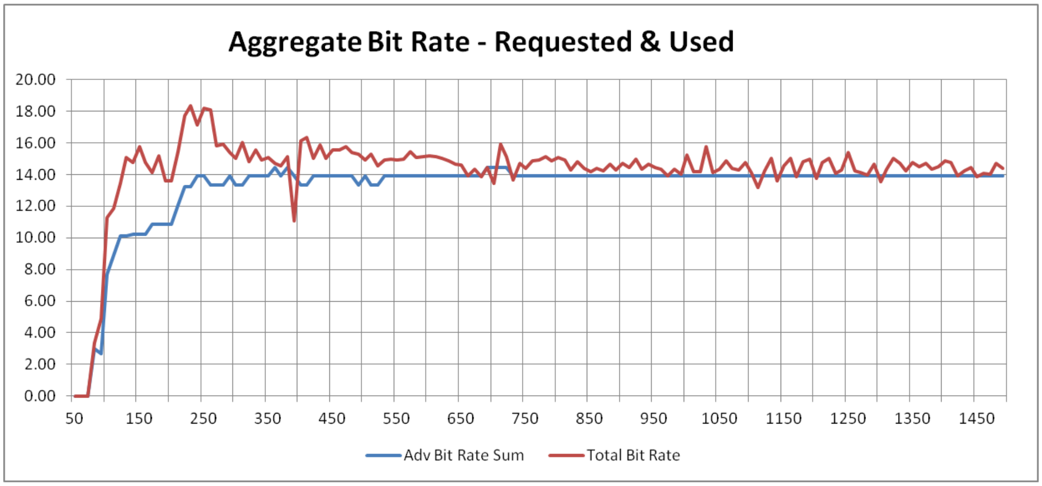

Figure 7a CMTS QoS, 12% Burst – Aggregate Rates

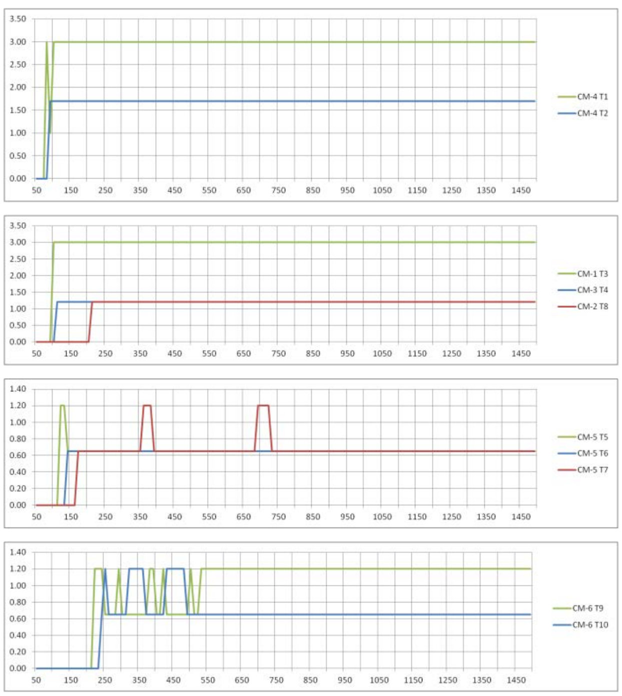

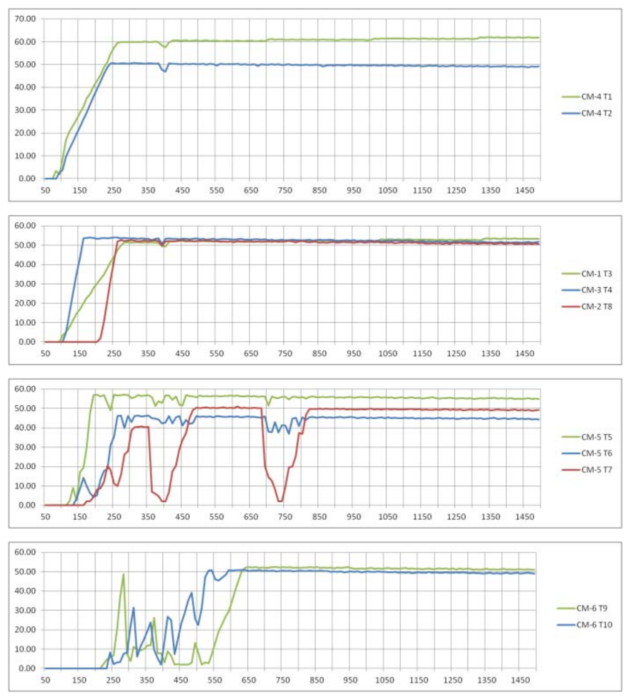

Fig 7b CMTS QoS, 12% Burst – Bit Rates

Fig 7c CMTS QoS 12% Burst – Buffer

CMTS QOS TEST RESULTS – CAVALRY TO THE RESCUE??

Before the CMTS QoS tests were even run, there were some high expectations that the CMTS could fix most all of the issues seen above for unmanaged ABR. With a dedicated Service Flow for every HLS client, it seemed like all clients would get their fair share of bandwidth thanks to the Weighted Fair Queuing scheduler in the downstream. With separate weighting on each Service Flow, even the HD clients should get a proportionately larger bandwidth share.

CMTS QoS Tests with 12% Burst Overhead

For the first CMTS QoS test, we intentionally kept the Service Flow Max Service Rates (MSR) just above the maximum video rate being offered. So VGA clients got a 2.5Mbps MSR for a 2.2Mbps max video rate and the HD clients got a 4.5Mbps MSR for a max 4Mbps video rate. The total aggregate MSR for the 10 Service Flows was 29Mbps, so it should come close to filling the channel. The results for this test are shown in Fig 7a-c.

Surprisingly, the aggregate bandwidth used peaked around 15Mbps and the average of the requested bit rates was ~14Mbps. So, this configuration had a network utilization of only 56%. This clearly shows that the HLS clients need a more substantive burst capability while in buffering mode. On the positive side, because the network was not congested, the system was extremely stable as can be viewed by Fig 7b & 7c.

Taking a closer look, both HD clients settle in quite comfortably at 3Mbps video rates with a 4.5Mbps MSR. For the VGA clients, there is a range of video rates from 4 @ 650Kbps, 3 @ 1.2Mbps and one @ 1.7Mbps. So, the CMTS scheduler was not able to correct all of the unfairness issues between the different VGA clients. Note that all but one of the clients settled at a bit rate that is less than half of its 2.5Mbps MSR.

The fact that network utilization was so low reinforces the complexity of the system with layers of flow control including TCP and the ABR clients. This area definitely deserves more research.

Note – video buffers in Figure 7c are generally quite stable (after startup) which bodes well for live Linear TV usage.

CMTS QoS Tests with Larger Burst Rates

The MSR was then opened up on the next set of CMTS QoS tests. The HD clients got a 9Mbps MSR and the VGA clients got a 4.5Mbps MSR, both slightly more than twice the max offered video rate. The key question now is the CMTS capable of fully utilizing the downstream channel while keeping the stability of the previous test. The results are shown in Fig 8a-c.

The network utilization does rise and appears to be close to the unmanaged ABR tests. The average bandwidth consumed after that initial startup is a tad less at ~23.1Mbps while the aggregate of all current bit rates selected averages slightly higher around 19.6Mbps after the initial startup. This bumps the network utilization up a hair to almost 80% of channel capacity.

Checking the stability of bit rates on Fig 8b shows that the two HD clients quickly snagged their max 4Mbps bit rate and kept it. Two of the early VGA starters locked onto a 2.2Mbps rate, but the remaining VGA clients did not fare so well. So, once a client can establish sufficient bandwidth demand, the CMTS QoS allows it to keep it. However, the remaining clients struggle with the remnants and show a lot of instability. This instability is apparent in the buffer depths of these clients in Fig 8c.

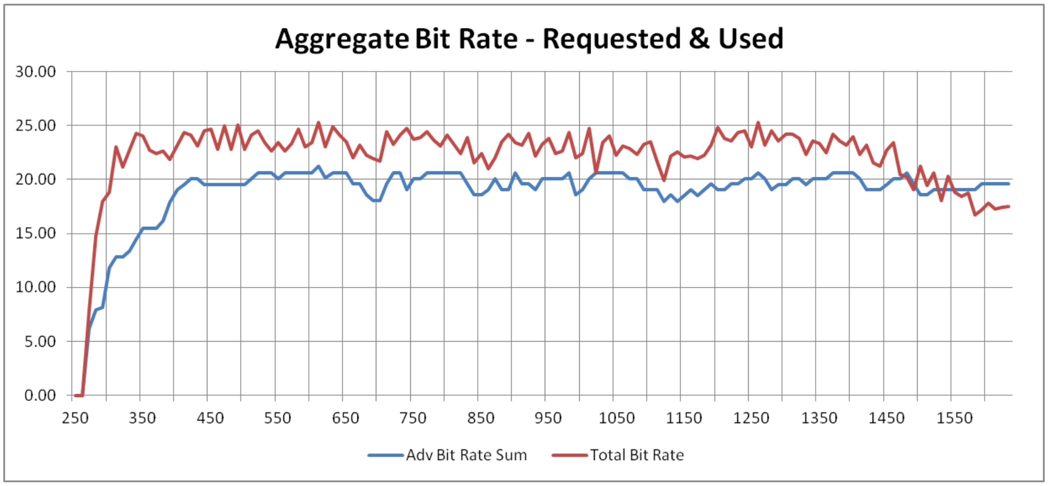

Figure 8a CMTS QoS, 100% Burst – Aggregate Rates

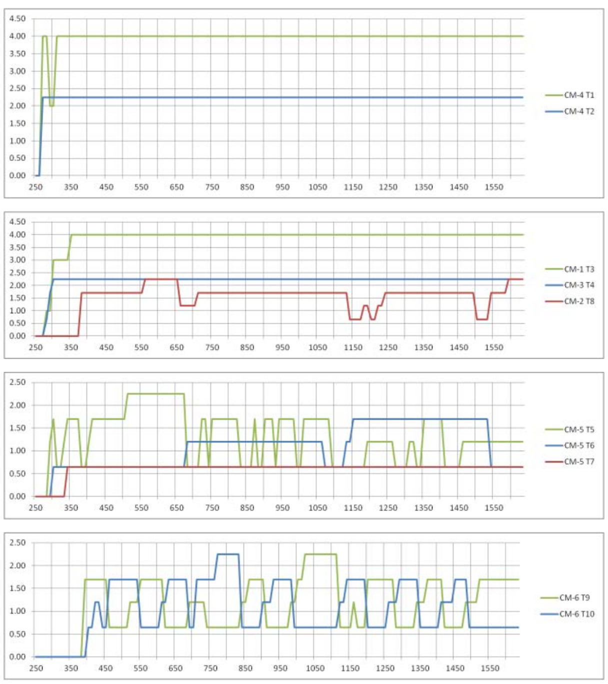

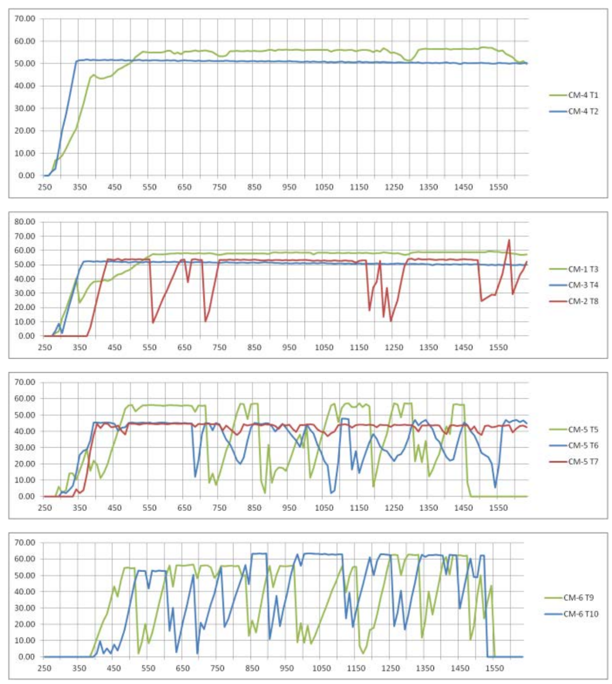

Fig 8b CMTS QoS, 100% Burst – Bit Rate

Fig 8c CMTS QoS, 100% Burst – Buffer

Another interesting observation is on the bottom chart for the two HLS clients sharing CM-6. These two clients are oscillating 180 degrees out of phase with each other. Most of the time the bit rates are bouncing between 650Kbps and 1.7Mbps. When increasing bit rates, there is a brief transition thru 1.2Mbps which increases the buffering overhead even more. It was noted that this cable modem was operating in the 2.4MHz WiFi band with multiple clients. This is another area for further investigation.

SABR SERVER TEST RESULTS – ORDER IS RESTORED TO UNIVERSE

In the previous CMTS QoS test configurations, the CMTS was trying to shape bandwidth to coax the clients into the optimal video rate selection. As was seen, the "drunken teenager" doesn't always cooperate. With the SABR Server testing, the bit rate choice is removed from the client and given to intelligence in the cloud. These results are shown in Fig 9a-c below.

The network utilization shows a definite improvement over the unmanaged ABR and CMTS QoS tests. While the average bandwidth consumed after that initial startup period is lower at ~21.8Mbps; the average video bit rate requested is ~10% higher at 21.1Mbps after the initial startup. This bumps the network utilization up over 84% of channel capacity. The SABR algorithms are still in its infancy, so there is a good opportunity to push the network utilization above 90% over time.

A first glance at the bit rates in Fig 9b might indicate significant instability. But in reality, the opposite is occurring as can be seen by the video buffers in Fig 9c. What is happening is that the SABR Server is intentionally adjusting the bit rates to try and maintain a constant video quality.

This is quite analogous to a Variable Bit Rate stream. The SABR Server must also trade off bandwidth needs between the different clients, so it performs the equivalent functions that a Statmux does in the MPEG world, only at a fraction of the computational complexity.

The net result of all of this is that SABR system produces a rock solid stable environment with excellent video quality and QoE. With the excellent control of the video buffer, this is ideally suited to distributing live Linear TV over ABR as well.

VIDEO QUALITY – THE REAL LITMUS TEST

As was seen in the previous unmanaged ABR and CMTS QoS tests, the steady state video bit rate was used to imply video quality. This metric no longer held with the SABR Server test configurations.

To do a more comprehensive study, our lab tests also measured the PSNR of every chunk delivered to each client. This allows a quantitative analysis of the delivered video quality. The appendix contains more details on how the PSNR is calculated.

Our analysis focus on two key aspects:

- *Percentage of chunks below 35dB*
- *Percentage of chunks above 40-45dB*

From our vast experience in video quality analytics, video below a PSNR of 35dB often translates to poor subjective video quality obvious to a common observer. As the PSNR improves much above 40dB, then the subjective video quality is getting past the point where a common observer might see any difference. In other words, having a PSNR that's too high means you are throwing bits away that a consumer will never observe.

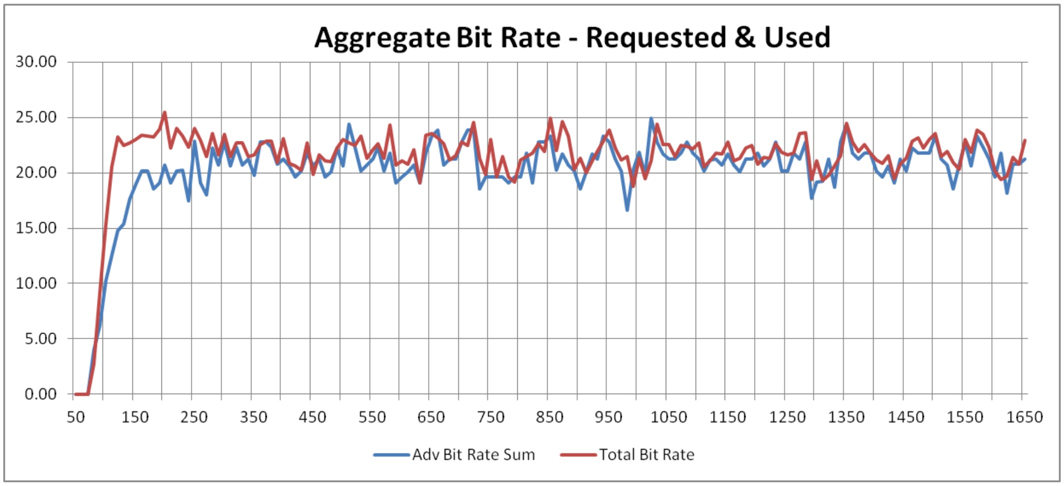

Figure 9a SABR Server – Aggregate Rates

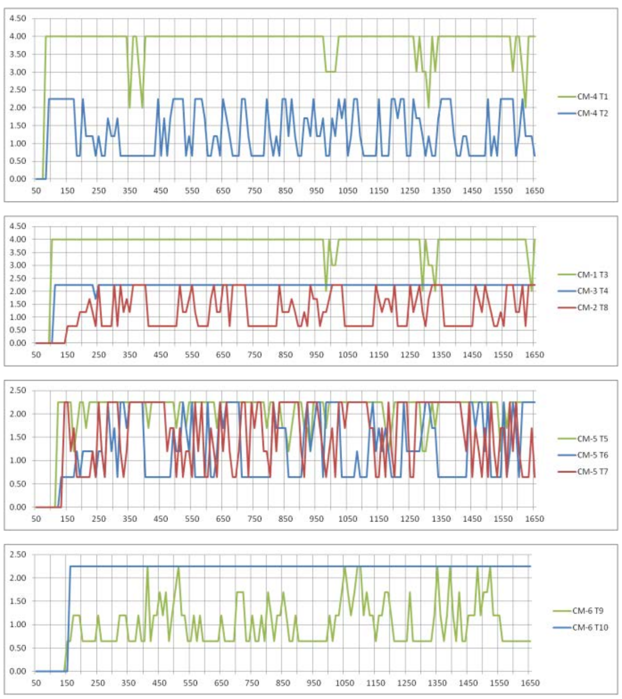

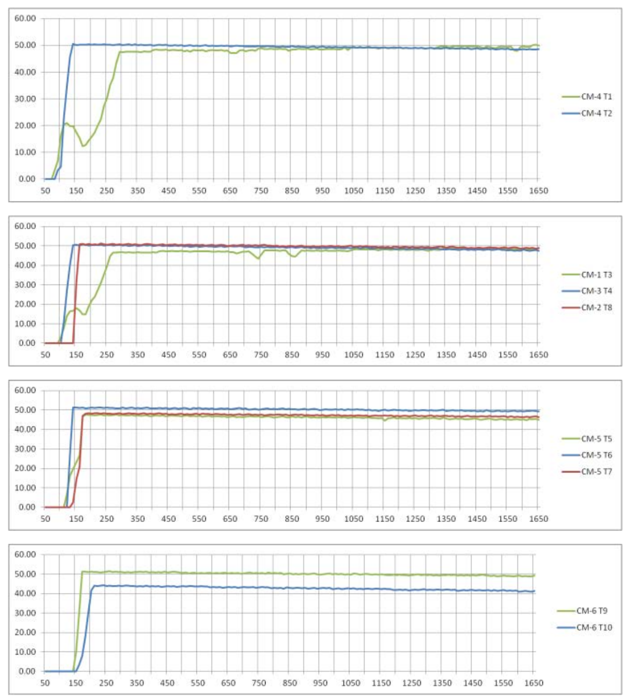

Figure 9b SABR Server – Bit Rates

Figure 9c SABR Server – Buffer Depth

The results of the PSNR calculations for all the previous tests are shown in Fig 10a-e. A separate PSNR curve is drawn for each and every client. These are a Cumulative Distribution Function (CDF) that shows the % of chunks on Y axis that match a given PSNR on X axis. Given these desired thresholds, the ideal curve would be a virtual step function that stays at 0% until 35dB and then reaches 100% by 40dB.

Figure 10a shows the PSNR results for the unmanaged ABR test. Five of the HLS clients have a significant portion (i.e. 25%-50%) of their chunks below the 35dB threshold. With no surprise, these map to the clients with unstable bit rates. The other five stable clients had 30%-70% of their chunks above 45dB, so lots of potential video bandwidth is being wasted. Also notice the spread (i.e. width) of the curves. This illustrates the unfairness in video quality between the different clients.

The results for the first CMTS QoS test with limited 12% burst rate are shown in Fig 10b. Remember that this test was very stable including the video buffer depth, but settled at a lower aggregate bit rate. However, the PSNR results were only marginally better than the unmanaged ABR case. Three clients had about 30% below 35dB PSNR while one horrible client had ~70% below 35dB. The others did extremely well, but in fact too good. Four clients had 70%-80% above 45dB, so lots of video bits going to waste. The PSNR spread between the clients is comparable to the unmanaged ABR test as different clients settled at different bit rates.

The next CMTS QoS test with full burst rate capability is shown in Fig 10c. The PSNR results of the 8 VGA clients are almost identical to the previous CMTS QoS test. The only notable PSNR difference was the improvement for the two HD clients. This is not surprising as both clients went from mostly a 3Mbps video stream to a 4Mbps video stream between these two tests.

Now let's take a look at the SABR Server PSNR results in Fig 10d. It was shown earlier that the video buffers were extremely stable while the video rates varied quite a bit. As can be seen, it produces significantly better PSNR results than both unmanaged ABR and CMTS QoS tests. Most of the clients have negligible >2% of chunks below 35dB, with one client at 12% and two clients at 18%. A closer look reveals that most of these chunks were actually within 1dB of the threshold. If a 34dB threshold is used, then seven clients drop virtually to zero, two clients at 2% and one client at 6%.

There are also significantly fewer chunks above 45dB with six clients at ~90% and the rest at ~80% by 45dB. The spread between the clients is dramatically reduced as well, with the spread between the best and worst being reduced by a factor of three. This is the ultimate litmus test that SABR provides video quality fairness across all clients.

The final figure 10e shows a previously unmentioned test configuration. This test limited the SABR Server to a max output of 15Mbps. The bit rate and video buffer charts were almost identical to the other SABR test, just scaled to a lower rate. However, the PSNR results for this test are very interesting. Because there is less total bandwidth, all video clients were given proportionately lower video rates; which in turn impacts the PSNR value.

Four of the clients had 30%-50% below 35dB, with the remaining clients less than 20% below 35dB. On the low end, this is very close to the results produced by the unmanaged ABR test. Put another way, SABR delivered comparable video quality in almost half the amount of bandwidth (i.e. 13.5Mbps vs. 23.6Mbps) while providing fairness and stability.

Another interesting comparison is with the first CMTS QoS test in Fig 10b. Both of these operated with just under 15Mbps of bandwidth so it is close to an apples to apples comparison. The SABR results show a significantly narrower PSNR spread of ~5-7dB while the CMTS QoS spread hovers around 15dB. This indicates SABR provides significantly better fairness between clients then the CMTS QoS approach.

These results show the real potential of SABR to increase the video capacity for operators while maintaining a very consistent video quality across _all_ of its consumers.

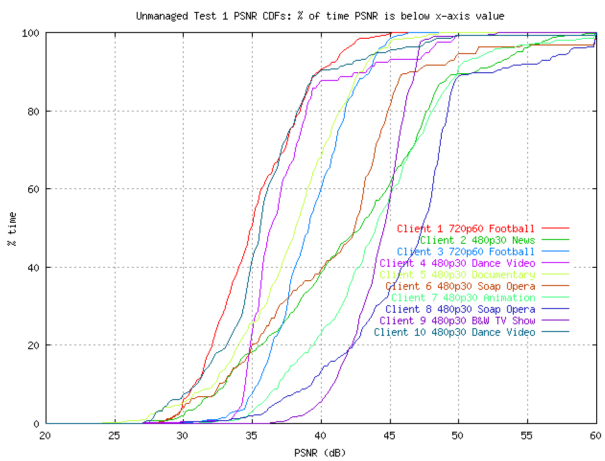

Figure 10a PSNR for Unmanaged ABR

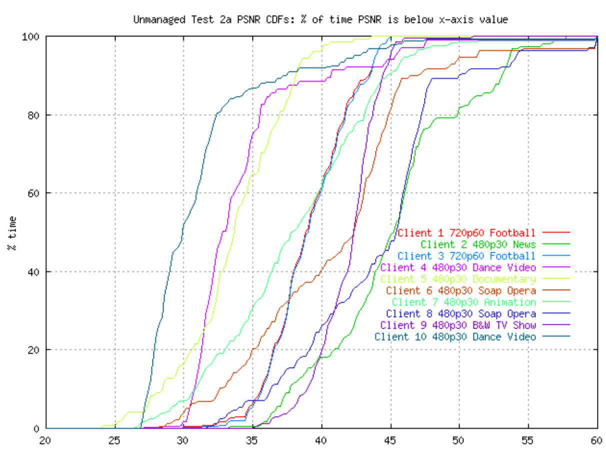

Figure 10b PSNR CMTS QoS, 12% Burst

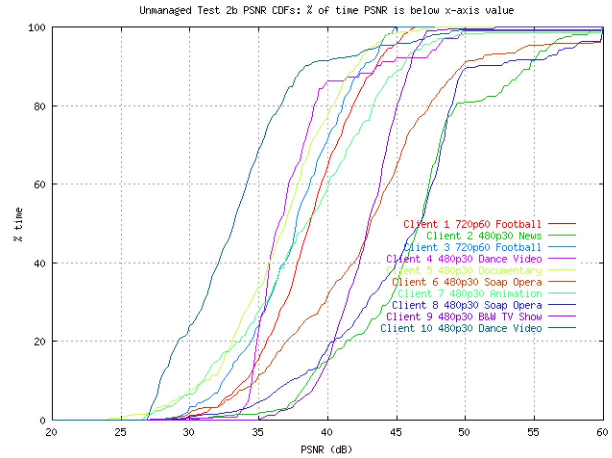

Figure 10c PSNR CMTS QoS, 100% Burst

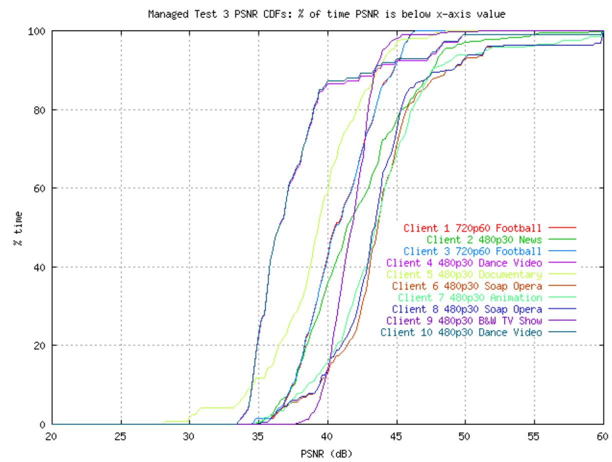

Figure 10d PSNR for SABR Server

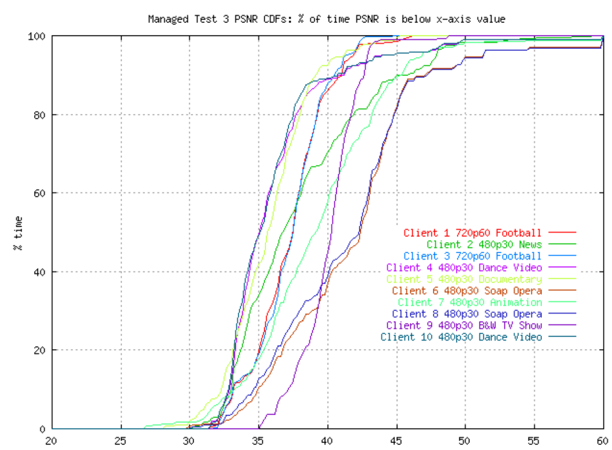

Figure 10e PSNR SABR Svr 15Mbps Cap

CONCLUSION

Operators have a number of challenges in offering a "managed" ABR video service. The limited bandwidth makes it challenging to support the demand of a large number of concurrent users while maintaining good video quality for each user. In addition, existing ABR client controlled distribution mechanism creates issues with QoE, network utilization, fairness and system stability. Our lab results presented confirms all of these conditions. In fact, fairness was an issue in multiple areas such as screen resolutions (HD vs. VGA), client software revision, shared home bandwidth and even the timing of when each client starts. While some clients were stable, others oscillated between bit rates and suffered video buffer under-runs.

A second series of tests were run leveraging enhanced CMTS QoS to try and fix the above issues. The results show only marginal improvements over the unmanaged ABR tests and the underlying issues of unfairness and instability still exists. Even had this approach fixed the problems, there are still concerns that the control plane could scale to support this solution.

Finally, our tests focused on a cloud based approach called Smart ABR (SABR) where the operator can regain control to provide a first rate video service with exceptional Quality of Experience while retaining the underlying benefits of Adaptive protocols. Our lab results show that all major shortcomings of unmanaged ABR can be addressed. The system becomes stable, network utilization improves and video quality fairness is provided across the entire client population while gracefully handling congestion and providing an improved user experience. Further tests suggest that SABR delivered comparable video quality in almost half the amount of bandwidth while providing fairness and stability

This lab testing is just the tip of the iceberg. Further research is warranted in areas such as optimizing SABR algorithms; mixing data traffic with the ABR video; investigating the impacts of WiFi in the home; further research into other CMTS QoS mechanisms; scaling testing to a much larger client population and understanding some of the subtle anomalies that appeared during our initial tests. It is important that the industry grasps the system dynamics for adaptive protocols.

With all of the promise seen with this initial testing, SABR should quickly become the future of IP Video.

ACKNOWLEDGEMENTS

In addition to our co-authors, I would also like to acknowledge Gary Hughes for his input and guidance into both the paper and our general research into adaptive protocols.

REFERENCES

[ADAMS] – "Optimizing Fairness of HTTP Adaptive Streaming in Cable Networks"; M. Adams, C. Philips, The Cable Show NCTA/SCTE Technical Sessions spring 2012

[CLOONAN] – "Competitive Analysis of Adaptive Video Streaming Implementations"; T. Cloonan Ph.D., Jim Allen Ph.D., The Cable Show NCTA/SCTE Technical Sessions spring 2011

[ULM] – "Adaptive Streaming – New Approaches for Cable IP Video Delivery"; J. Ulm, T. du Breuil, G. Hughes, S. McCarthy, The Cable Show NCTA/SCTE Technical Sessions spring 2010

[WHITE] – "Reclaiming Control of the Network from ABR Video Clients"; G. White, J. Ulm, The Cable Show NCTA/SCTE Technical Sessions 2012

APPENDIX – TEST SETUP DETAILS

Table 1 shows the details of the iOS client configuration and the media streams retrieved by the clients. Two iPads were connected to CM-4; one retrieving the complex 720p30 Football video and the other retrieving the less complex 480p30 News clip loop. Single iPads were connected to CM-3 and CM-1 with CM-3 delivering the complex 480p30 Dance Video clip and CM-1 streaming the complex 720p30 Football sequence. CM-5 had three iPads associated to it; each receiving 480p30 sequences with one of high video complexity and two of lower complexity. One iPod was associated to 2.4 GHz CM-2 and retrieved a low complexity 480p30 Soap Opera sequence while two iPods were associated to 2.4 GHz CM-6 with one iPod getting a less complex 480p30 black-and-white (B&W) TV clip and the other pulling the complex 480p30 Dance Video sequence.

Test configuration 1 had the goal of determining the baseline performance of unmanaged HLS clients retrieving the various media content over a DS BE SF as managed by the BSR2000 CMTS for each CM. The DS SF Maximum Sustained Traffic Rate (MSR) was set to 12 Mbps, the Maximum Traffic Burst (MTB) to 8 kBytes, and the Minimum Reserved Traffic Rate (MRR) to 0 Mbps.

Test configurations 2a and 2b entailed configuring per-HLS stream SFs so that clients retrieving the 720p30 content might be differentiated from those retrieving 480p30 content. To that end, test configuration 2a entailed applying SF MSR of 2.5 Mbps to the 480p30 content, which had maximum bit rate variant of 2.2 Mbps, and MSR of 4.5 Mbps to the 720p30 content which had maximum bit rate of 4 Mbps. Test configuration 2b doubled these MSR values to 5.0 Mbps for 480p30 and 9.0 Mbps for 720p30 clients. As before, MRR was set to 0 Mbps and MTB to 8 KBytes.

The tests were conducted by launching the iOS client Safari browser and pointing it to the URL of the desired media's variant playlist on the unmanaged HLS server connected behind the CMTS.

The PSNR is calculated as follows. Decoded video was subtracted from the uncompressed video. This difference provided the compression noise (distortion) present in each pixel of the video at that particular rate. Mean Square Error (MSE) was calculated by squaring the difference and averaging over the picture. As the maximum luminance value is 255, PSNR was calculated by taking the ratio of square of 255 and MSE. This was converted into dB by taking log of PSNR and multiplying it by 10. As this measures the compression noise, it provides a rough measure of visual quality. Higher is the PSNR, lower is the compression noise and roughly better visual quality.

Client No.	iOS Device Type	iOS Model	iOS Version	Cable Modem	SBG6580 WiFi Channel	ABR Media Stream	Video Complexity & Stream Rates (Mbps)
1	iPad 4	MD513LL	6.1.2	CM-4	157	Football 720p30	High @ 1,2,3,4
2	iPad 2	MC769LL	6.0.1	CM-4	157	News 480p30	Low @ 0.65, 1.2,1.7,2.2
3	iPad 2	MC705LL	6.0.1	CM-1	48	Football 720p30	High @ 1,2,3,4
4	iPad 1	MB292LL	5.0.1	CM-3	153	Dance Video 480p30	High @ 1,2,3,4
5	iPad 1	MB292LL	5.1.1	CM-5	36	Documentary 480p30	High @ 1,2,3,4
6	iPad 2	MC769LL	6.1.3	CM-5	36	Soap Opera 480p30	Low @ 0.65, 1.2,1.7,2.2
7	iPad 2	MC769LL	6.0.1	CM-5	36	Animation 480p30	Low @ 0.65, 1.2,1.7,2.2
8	iPod	MC540LL	5.1.1	CM-2	6	Soap Opera 480p30	Low @ 0.65, 1.2,1.7,2.2
9	iPod	MC540LL	5.0.1	CM-6	1	B&W TV Show 480p30	Low @ 0.65, 1.2,1.7,2.2
10	iPod	MC540LL	5.0.1	CM-6	1	Dance Video 480p30	High @ 1,2,3,4

Table 1. Adaptive Bit Rate HLS Client & Video Stream Details

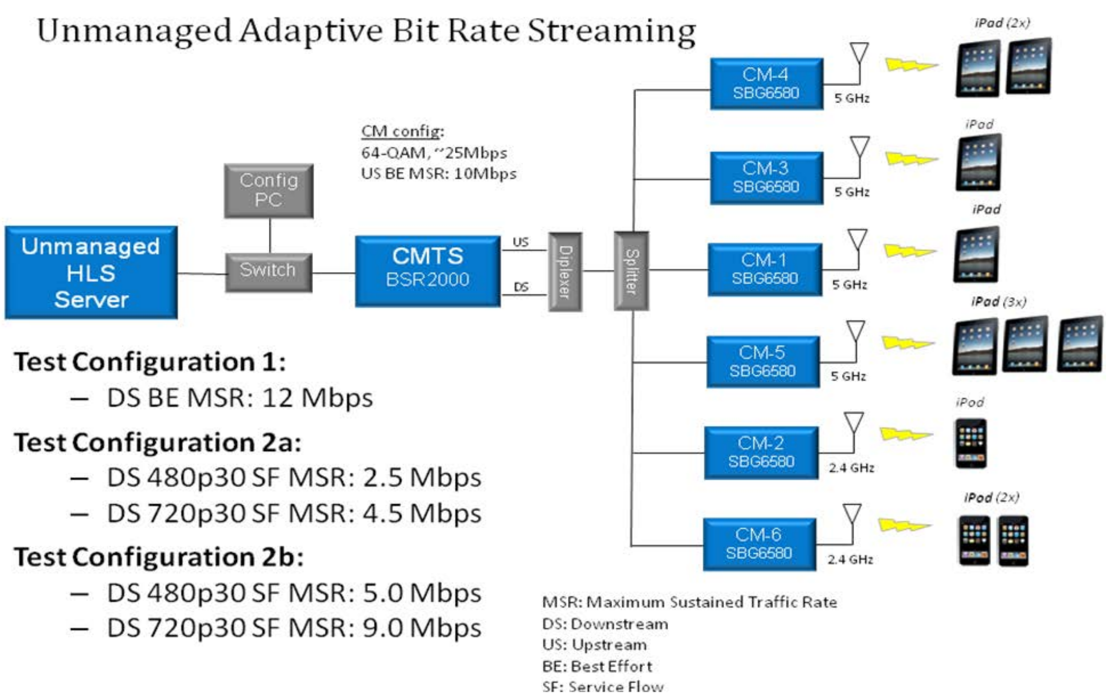

Figure 5 Adaptive Bit Rate Streaming Test Setup

SOFTWARE DEFINED NETWORKING AND CLOUD – ENABLING GREATER FLEXIBILITY FOR CABLE OPERATORS

David Lively
Cisco

Abstract

Software-defined networking (SDN) promises tremendous flexibility for operators, especially as their application environments become more dynamic with cloud computing. As more content is becoming available to consumers via IP every day – from sources both internal and external – managing traffic flows for content and for the associated bandwidth across both the regional network is becoming more critical. Gaining real-time access to network analytics allows the network to optimize routing when connecting clients to content. This paper focuses on the content delivery use case, and discusses how cloud computing and SDN can work independently and together to deliver the highest quality experience to their customers.

CLOUD AND SDN – DIFFERENT BUT THE SAME

Cloud computing has delivered tremendous flexibility for application developers. The general definition of "cloud" is resources, abstracted from the physical hardware, delivered elastically and on demand. Because of this abstraction, applications can be deployed and scaled without the need to physically configure and deploy new hardware. Initially, this flexibility was primarily under the control of the administrators, and presented to users through portals to request services. Traditional and legacy applications could benefit from this newer, more flexible deployment model, but the real breakthroughs started to come when new applications were developed to call APIs exposed by the cloud platform directly. Subsequently, intelligence was written into the application itself to call APIs from the cloud platform to allocate additional resources when the application needed, and to turn them off when they were no longer needed. The cloud model initially applied mostly to compute, with the virtual machine, but has since expanded to include network and multiple different storage models (object storage, volume storage, etc.), as well as higher level services such as database services and security. These services further abstract the physical infrastructure beyond compute and storage to services that the application can directly use.

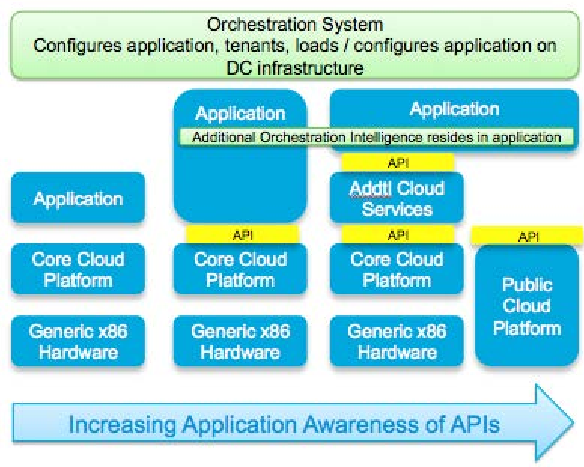

This model of higher layer services abstracted from the physical infrastructure starts to more closely resemble software-defined networking. Constructs have existed in networking for some time to allow a single physical router or switch to be segmented into multiple virtual and logical routers or switches – VLANs, VRFs, etc. SDN takes this concept one step further by providing the ability to define and request resources across a larger scope of the network. As applications became decoupled from the physical infrastructure and were able to move, grow, and contract as needed, the network connectivity between them needed to move, grow, and contract accordingly, and automatically, with the application. A heavily virtualized application world, coupled with increasingly mobile users, means that we don't know when the network flows are going to come in or where they're going to come from. Leveraging business intelligence to help drive network behavior through SDN can help the network deliver the flexibility and adaptability to handle these new dynamic flows.

GAINING THE MOST FROM SDN WITH BETTER NETWORK INTELLIGENCE

Consumer devices are accessing content from the network from both fixed and mobile locations. As fixed and wireless broadband speeds increase, this increase in the number and range of devices, along with the average bandwidth of the content flows, is having a dramatic affect on the network and its available bandwidth. Software-defined network can potentially help optimize flows on the network, but only if the SDN controller has access to information about the state of the network. Traditionally, this data has been collected by the network and stored for later analysis, with access to information delayed on the order of minutes, or even hours / days depending on the analysis being done. While access to historical data can be tremendously beneficial for future network planning, the ability to analyze the data in real time and make decisions on that analysis can have a dramatic affect on the performance of bandwidth or delay-sensitive applications such as video. For example, many operators are using caching techniques with CDNs to minimize the impact of the increasing content streaming on their networks, but as more and more users are accessing content from a greater range of locations, network intelligence can be used to further optimize those requests.

POTENTIAL SDN USE CASE – OPTIMIZED REQUEST ROUTING IN DYNAMIC CDNS

Routing content requests to the cache that optimizes the end-user's experience while at the same time minimizing the impact on the network requires consideration of multiple factors from the physical layers of the network up through to the application layer. In addition to determining which caches have available capacity for handling incoming requests, the state of the network can be utilized to optimize the experience.

Determining the optimal cache to source content from

When choosing which cache to route a content request to, the network and application must consider more than just the "closest" cache from a network or hop count perspective. When multiple caches are available over disparate paths, the closest cache may be sub-optimal from a network standpoint for delivering the video. Latency and jitter can play a large role in the end user's viewing experience, causing buffering issues, etc. Available bandwidth is another consideration. A network link with low latency and a shorter routing path, metrics normally considered by dynamic routing protocols, might still provide an unsatisfactory viewing experience for the end user, especially for single bitrate streams that are not able to adjust to congestion conditions.

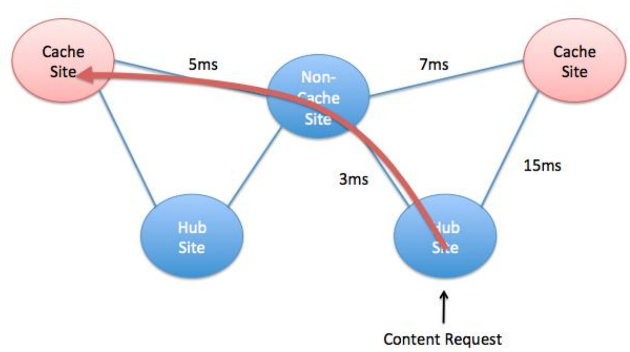

Latency-Optimized Cache Choice

Additionally, if the initial cache request is known or assumed to be part of a longer duration flow, that can also be taken into account when routing requests to the optimal cache. The SDN controller could ask the routers to add timestamp information to packets and start building a database of latency information between different nodes, and along different paths between the same two nodes. Once the operator has obtained quantitative analytics, such as latency, from the network, they now have the ability to define policies that leverage that business-level intelligence to choose the optimal location for routing content requests, and ultimately to program the network itself in response.

<u>SDN Options for Content Requests</u>

In addition to using network intelligence to determine the optimal cache to route requests to, we also have the option with SDN to specify and optimize the routing of specific flows to be different from what typical routing protocols would specify. We will discuss two potential approaches for using SDN technologies in this CDN use case. The first involves "forcing" the content flow to take a different path to the cache than would be chosen by the network using standard routing metrics. The second approach focuses on optimizing the primary path for the content flow by affecting the routing or handling of other traffic on the primary path. These two scenarios can also apply outside of CDN use cases, and can be used for any scenario where multiple paths exist to deliver content, and where the path that is most optimal from a routing protocol perspective is not the most optimal for the viewing experience.

<u>Modifying the path for a Specific Content Flow</u>

There are many conventional methods that network engineers use today to modify the path of particular flows through the network based on specific policies. Engineers can influence next-hop behavior for packets, can use traffic engineering to define specific paths through the network, can used QoS marking to specify specific output queues for traffic, etc. However, all of these methods typically require pre-configuring specific features on the devices in the network and won't scale to handle large numbers of per-flow requests on demand. SDN provides us with a way to programmatically have the network move the flow to a different path through the network, and then remove those routes or override metrics once the flow is done.

In this use case, we are going to assume that the primary path to the cache has become congested, or is undesirable for other reasons, and a secondary path exists that meets the required policies exists. These policies could be network-oriented (such as bandwidth requirements, delay requirements, etc.), or could even be business-oriented policies. Perhaps some links and routes are being leased from other carriers or have more expensive peering costs, so even financial decisions can be used to influence traffic patterns.

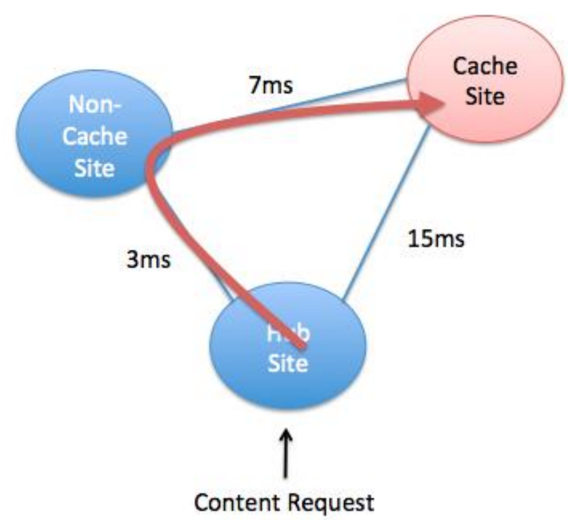

Latency-Optimized Path Choice

Thus, for a specific flow or class of flows, we want to have the routers send the traffic via an alternate path. We do not want to alter the path for any of the other traffic, so we can't change the general routing metrics for the links overall for all traffic. The first step is to identify a particular flow, or class of flows, and there are multiple mechanisms that exist today to do this. In the most trivial scenario, a flow could simply be identified by source and destination IP address. In practice in cable networks, many end-points are behind NAT devices such as home routers, and many end points may look like a single IP end point. In this case, the operator can get more granular and identify a flow by source and destination IP, plus flow type as inferred by the source and destination TCP / UDP ports. These are both readily doable using information in the packet header itself, and routers can generally make routing decisions in hardware using information from the standard header. You can also get even more granular and use deep packet inspection techniques to identify a particular flow.

Once you've identified a particular flow, you need to process packets in that flow according to the technique being used to modify the routing of the packets. This could be done by encapsulating the packets with new header information and placing them into a specific tunnel across the network using GRE or MPLS. They could also be "tagged" with a new service header or identifier in the packet header that is specific to the service. Minimally they could simply be assigned to a policy that governs their next-hop routing interface.

One method to leverage the programmatic capabilities of SDN would be to use the controller to dynamically program ACLs into the routers, which would be used to map packets to the appropriate policy. An external SDN controller would have the broader view of the overall network, including where the caches are located, and what various network metrics such as latency, delay, bandwidth, etc. look like on each segment. When a request for content is made, the controller can determine the optimal route to the selected cache.

Another method would be to use the service header described above and have routers assign packets to a particular path based on the service header. These "service paths" through the network can be programmed into the routers from the SDN controller. Since the service paths only need to exist for the duration of the flow, they can be added and removed as needed by the SDN controller.

Optimizing the Primary Path by Dealing with Other Traffic

Similarly to moving a single flow or class of flows to an alternate path, it's possible to optimize the "favored" path for a flow by using SDN to modify the forwarding and QoS behavior for other types of traffic. The simplest method is to simply mark the priority flows such that they are processed by higher priority queues on a device. However, in some serious congestion scenarios it may not be desirable to simply drop lower priority traffic, but instead re-route that traffic over a secondary path. In this case you could use the same technique described above for a single content flow, but use it instead to classify "other" traffic that is less bandwidth and delay sensitive (or a lower priority for any number of other reasons) and to modify the routing metrics or path selection for that overall class of flows. In this way, the overall traffic on the link goes down, providing a better experience for flows that need to stay on the link.

Other Factors – Path to Origin Server

In a caching scenario, it's not always the path from the user to the cache that needs to be optimized, but the path from the cache to the content origin server. In the case that high amounts of new content suddenly need to be cached (such as a new viral video or other event), the paths from the origin servers to the caches may become the bottlenecks for performance. A cache (cache 'A') that looked ideal when evaluating the network performance and load between the cache and the user may no longer be optimal if all of the content needs to be source from the origin server, and the path between the origin server and cache 'A' is highly congested. In this case, cache 'B' might be a better cache if the path from it to the origin server is better able to handle the load, and if the path from cache 'B' to the end user can support the required policies.

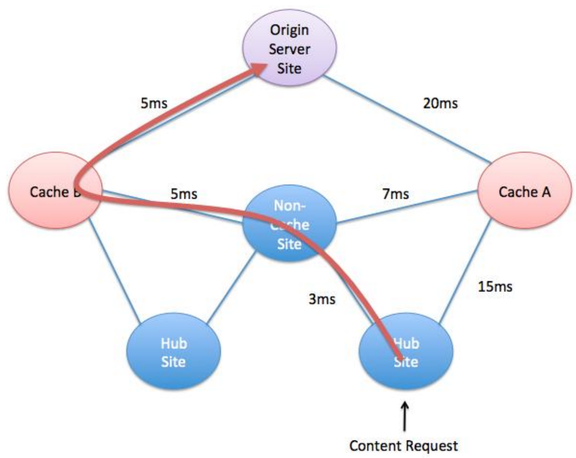

Latency-Optimized Cache Choice Considering Path to Origin Server

SDN for Modifying Flows in CDNs

As discussed at the beginning of this section, many approaches already exist today for both classifying packets, and modifying the routing of those packets through the network on a per-flow basis. The various approaches use existing routing protocols, policies, tunneling technologies, etc. However, most, if not all, of these approaches would be impossible to do dynamically, on a per-flow basis (as requests are made) without having both access to real time network analytics to make the decision as well as programmatic interfaces to the routers to apply the changes. SDN provides us with a way to have a controller with a broader view of the network, including non-network analytics such as business policies that can access devices in the network through programmatic interfaces to modify packet routing behavior.

LEVERAGING CLOUD COMPUTING FOR DYNAMIC APPLICATIONS

While networks are being enhanced to take more advantage of the dynamic capabilities of programmatic interfaces with the network, and enhanced analytics, the applications themselves are becoming more dynamic and programmed to take advantage of this additional flexibility in the network. By working together, applications can take advantage not only of the cloud computing capabilities of spinning up additional application instances on demand, but also of the intelligence available in the network to spin up resources in the best location when multiple are available.

Caching Example – Using Cloud for the Spikes

Referring back to the CDN example discussed in the CDN section, one of the first things that the CDN could check is whether or not a cache has enough capacity to support servicing incoming content requests. Caches are typically designed and engineered to handle typical loads on the network plus a predicted amount of burst capacity, based on analyzing historical data such as number of consumers in a given area, type consumer viewing behavior, busy hour requests, etc. While this will handle the vast majority of requests for content, there exist both seasonal "expected" dramatic increases in content viewing (such as around major sporting events like the Super Bowl or the World Cup) as well as dramatic unexpected spikes in viewing for things like suddenly viral videos.

If the network operator has built enough caching capacity to handle these short term, dramatically spiked increases in capacity, much of the caching capacity would sit idle and wasted the rest of the time. Since many of the caching servers are now able to be run as virtual machines on standard x86 hardware, cloud computing technologies can be used to "spin up" additional caches for both expected and unexpected events, and can then be brought back down to allow the compute resources to be used for other applications that see spikes in demand at different times, or whose usage is not as time-critical as video delivery. This capacity can also be brought up at locations where there is excess bandwidth to handle the additional load being placed upon the network. Operators can both maximize the number of content requests that can be handled by caches, and also minimize the impact on the network by locating those caches in the most strategic locations from a network perspective.

As additional caching capacity is brought online, it can dynamically be added to controller that is assigning content requests to caches. Additionally, the additional capacity is made known to the SDN controller(s) which can then take those caches into account (as well as their locations in the network) when determining how to deal with new flows.

Geographic location can also be used as a part of the decision process. For example, if there is an unexpected spike in content requests during prime time television viewing on the east coast of the United States due to a promotion, additional capacity could be proactively brought online in the central and west coast regions to handle the anticipated increased demand. Additionally, content requests could be sent to the new and existing caches to allow them to start caching content ahead of anticipated demand. If the demand did not materialize as expected, the additional caches that were started up can be shut down, with their compute resources being returned the available pool, and existing caches will simply age out the content as other requests come in.

LEVERAGING CLOUD COMPUTING APPLICATION DESIGN PRINICPALS IN NETWORKING

Cloud computing capabilities are helping to bring about complementary capabilities in networking. Cloud computing is further enhanced when it starts taking advantage of the new APIs available from the network, and using network location and state information in its decision-making process on when and where to spin up new capacity for applications. Additionally, the software architecture design principals being used by new applications being developed specifically for the cloud can provide additional hints on how to look at network design for the cloud. Several design principals are emerging, including:

- Fault-tolerant design that is adaptive and programmable through APIs
- Dynamic instantiation of new services and capabilities for the application
- Scaling out (adding additional nodes for a distributed function) vs. scaling up (adding compute or storage capacity to individual nodes)
- Heavy focus on automation

The new application design principals are both influencing network designs (including SDN) as well as being designed to take advantage of the more programmable nature and design of today's networks.

Fault-Tolerant Designs

Cloud applications are being architected such that the application will continue to run even in the event of physical infrastructure failures. Failures of servers or server components, networking gear, and even physical data centers won't bring down well-designed cloud applications. The applications are designed using multiple loosely coupled components that are distributed across multiple different availability zones. Typically each component can be loaded dynamically as needed, and scaled up or down as needed. All of this is starting to translate to network design, as networks need to be able to accommodate application communication across multiple servers, racks, and entire data centers. When loads shift due to outages, networks need to be able to quickly accommodate new flows, and to rapidly apply changes in network policy.

Dynamic Instantiation of New Services

Many applications need to leverage network-based load balancing to help scale individual application components and distribute loads across availability zones. As the application scales out, additional load balancers need to be dynamically instantiated and added into the architecture to handle the new node addresses. As load drops, load balancers can be taken offline, with their resources freed up for different applications or services.

Newer video applications being delivered from the cloud can also take advantage of the dynamic instantiation of services. For example, when a user requests a particular piece of content, that content may need to pass through a just-in-time transcoding or encryption application. As these applications are developed to be loaded dynamically in a cloud computing platform, a base amount of capacity can be kept online, with additional capacity being brought online in "standby" mode as loads increase. Understanding of network analytics is important for network-based services such as these that require significant bandwidth, and can greatly benefit from intelligent placement or routing due to the bandwidth demands.

Scaling Out vs. Scaling up

The historical approach to handle capacity growth, called vertical scale or scaling up, involves turning up an application on bigger and faster hardware to handle a larger load. Modern cloud applications are being developed to handle increased loads by scaling out - adding additional instances of the application on additional, smaller compute nodes as load increases and balancing between them. Networks are now often being designed in similar fashion, with smaller fault-domains per switch pair, and scaling through the addition of more switch-pairs. In this fashion, network capacity can be added as compute and storage capacity is added without increasing the risk of a single failure taking out mission critical applications.

Heavy Focus on Automation

Finally, applications would not be able to scale dynamically without a heavy focus on automating every step of the process. Manually adding or removing capacity, or bringing up new services for the application can take significant time, and has significantly more potential for errors to be introduced into the process through manual data entry. This applies to all aspects of the application and its deployment, from the compute to the storage to the network connectivity, network services, and ultimately application configuration itself. The other half of the automation focus is monitoring. Without the ability to accurately monitor and measure application and network loads, it would be impossible to know when to scale up or down capacity. Programmatic APIs make both monitoring and automation possible. Application designers, as well as network designers, set the business policies and rules up front, and then let the automation systems take care of things from there.

CLOUD AND SDN – WORKING TOGETHER FOR GREATER CONTROL AND FLEXIBILITY

Advances in cloud compute brought about significant improvements in both the speed and flexibility of application deployment by abstracting the resources used by the application from the physical infrastructure providing those resources. This abstraction helped bring about the ability to dynamically bring up and down resources as needed, without the need to configure physical hardware. As developers learned how to use this flexibility, software architectures started evolving to take advantage of the new capabilities. This new generation of applications and architecture thinking helped bring out similar changes in network applications and design. The introduction of programmatic interfaces on network devices enables software to define network behavior, utilizing business intelligence and policies in lieu of just network metrics. In some cases the intelligence resides within the application itself, and the application is able to program and control the network directly. In other cases the intelligence resides within a specialized SDN controller that acts as an abstraction between the application and the network.

However, before you can take full advantage of the network abstractions and programmatic capabilities available with SDN, you need to start with monitoring and visibility of the network – analytics. Once you have a view of what the historic and real-time state of the network is, and you understand the business rules required by the application, you can use the SDN controller to enact custom routing on the network.

It starts with setting policy in the controller, which communicates the policy down to the network devices who enforce it by applying new forwarding rules, encapsulating packets into tunnels, etc. As the application developers start to understand the network analytics and abstractions available, they can use business rules to define policies in the controller and let the controller define network routing based on those policies.

While SDN adds flexibility, the use cases we have talked about in this paper are more about using SDN to manage the exceptions rather than replacing standard network protocols and router control planes altogether. The network is already highly adept at routing and moving packets based on network state. But when applications need higher-level business intelligence to make a better decision for the application, SDN can be used to program the exceptions that need specific or different handling. Exception-based SDN use cases allow the network to continue to do what it does best, and allow the SDN controller (and ultimately the application) to define behavior for specific flows, such as content flows within CDNs.

The Power of DOCSIS 3.1 Downstream Profiles

John T. Chapman, CTO Cable Access BU & Cisco Fellow
jchapman@cisco.com
Cisco

Abstract

DOCSIS 3.1 will offer a more robust PHY that supports technologies such as OFDM and LDPC. With these technologies comes the opportunity to customize modulation profiles for groups of CMs to take advantage of the variation in SNR on an HFC plant.

This paper will explore the use of the modulation profiles in the downstream path.

TABLE OF CONTENTS

Table of Contents 1
Introduction 2
  Disclaimer 2
  OFDM 2
  LDPC 2
Defining The Need 3
  Is One Profile Good Enough? 3
    Axiom 1 3
    Axiom 2 3
  The OFDM Paradox 3
  New Spectrum Opportunities 4
    Below cutoff: 4
    Above cutoff: 4
Quantifying The Value 4
  Channel Study 4
  Efficiency Analysis 5
Building The Mechanism 6
  Codeword Builder 6
  Latency 6
System Operation 7
  Details 7
    Profiles and OFDM Channels 7
    Profiles and Paths 7
    Multiple Paths 8
    Path Assignment 8
    Service Flows 8
    Bonding 8
  Bandwidth Management 9
    MAC or PHY? 9
    Channel Capacity 9
    Time-Varying Channel Capacity 9
    Frequency-Varying Channel Capacity 10
  Profile Management 10
    Static Profiles 10
    Dynamic Profiles 10
    Measure and Sort 11
Choosing the Number of Profiles 12
  HFC Plant Variations 12
    MMM and Data 12
    Penalty Box 12
    Modulation 12
    Geographical Differences 13
  IP Multicast 13
  Profile Update 14
    Whole Profile Update 14
    Migrate to a New Profile 14
  Profile Changing 14
    Pause / Play 14
    Add / Drop 14
  Probing 14
    Predict 15
    Test & Measure 15
  The 4+1 Proposal 15
Summary 15
References 16

INTRODUCTION

Disclaimer

The ideas described in this paper are being considered to become part of DOCSIS 3.1 specifications.

DOCSIS 3.1 specifications are still under development. The following represents the author's current thoughts on what downstream profiles in DOCSIS 3.1 might look like, and do not represent actual decisions made regarding the final form of the specifications or technology.

Anything could change. Seriously.

OFDM

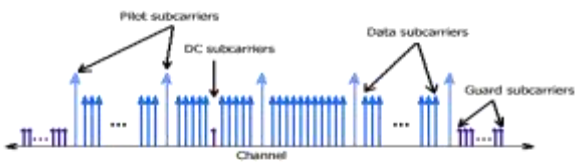

Figure 1: OFDM

As of this writing, DOCSIS 3.1 – the next version of DOCSIS – is being defined. While the ink is still drying and the mystery of what DOCSIS 3.1 will hold has yet to be revealed, one of the more certain outcomes is that DOCSIS 3.1 will move from relying on the tried and true single-carrier QAM (SC-QAM) technology for the PHY, and will incorporate OFDM – Orthogonal Frequency Division Multiplexing.

OFDM is really just a collection of very narrow QAM subcarriers as shown in Figure 1. Because they are orthogonal, these narrow QAM subcarriers can be placed very close to each other or even overlap each other, making for much more efficient use of the available spectrum.

For example, an OFDM channel might be as large as 192 MHz and contain 3840 QAM subcarriers, spaced every 50 kHz, or even 7680 carriers, spaced every 25 kHz.

OFDM allows each subcarrier to have its own modulation value and amplitude (or none at all in the case of a muted subcarrier). It is this interesting property that allows OFDM to fit a downstream transmission path like a glove.

The description of the modulation type for each subcarrier and its amplitude value is stored in a data structure simply referred to as a profile.

Once you can tailor such a glove, just how many gloves, or profiles, are needed for a downstream HFC plant? How are these profiles managed?

LDPC

A second part of the DOCSIS 3.1 story that is interesting is the adoption of LDPC (Low Density Parity Check). LDPC is a FEC (Forward Error Correction) technique that can be be used to correct bit errors due to noise. LDPC is more powerful than its DOCSIS predecessor Reed-Solomon. As such, LDPC is a major factor in allowing higher order modulation to be used.

So while LDPC allow a great density of bits per hertz, OFDM allows for a better customization of the frequency spectrum and noise mitigation.

DEFINING THE NEED

Is One Profile Good Enough?

The answer to this question is not obvious. In fact, there are proponents on both sides of this debate.

The argument for one profile would be to pick a modulation value that is good enough. For example, today's plant is almost entirely 256-QAM. There is rarely a need to do anything more or less. Some operators will put 64-QAM in the roll-off regions of the HFC plant, but that is more of an exception rather than a rule.

Thus, in a HFC plant with OFDM and an improved FEC, the new norm could be 1024-QAM. To enforce that, any part of the plant that could not support 1024-QAM with the new CMs would not be upgraded until it did support it. In such a scenario, no more – and no less – would be needed.

The other side of the argument is that either the plant cannot or is not always upgraded everywhere, or that even with a well-maintained plant, there is enough natural variation is SNR (signal-to-noise ratio) that extra performance can be squeezed out.

Axiom 1

It is easier to downgrade than upgrade.

If some of the CMs in the HFC plant are not working well, either due to interference by noise-like carriers such as digital TV ingress, or due to an insufficient SNR margin, the profile can be modified to accommodate and thus alleviate the problem. This action would result in a profile that has lower throughput but higher reliability.

But what if the system is already working? Hypothetically, the profile can be upgraded as well. In practice, if upgrading a profile is not done right, then CMs may drop off. This improvement often flies in the face of "if it ain't broke, don't fix it."

The likely outcome is channel degradation over time that is hard to recover from.

The take-away is that if profiles are to be updated, then a robust system for measuring, testing, and deploying profiles is needed. Such a system would benefit from the ability to create additional profiles that can be tested before deployment.

Axiom 2

The decision to sort is easier than the decision to deny or downgrade.

For a single profile system, when a CM is in a bad part of the HFC plant, the choice is to downgrade everyone or deny service to that particular CM. When modifying an OFDM profile in a system with only one profile, <u>all</u> CMs are impacted by definition. So, if one CM is having problems, scenarios could exist where fixing the service for one CM may break the service for other CMs.

With a multi-profile system, unlucky CMs can be moved to a lower throughput profile without disturbing the rest of the CMs. In a system with multiple profiles, CMs that are in trouble can be moved to isolated profiles, or to profiles that have a lower throughput and higher robustness.

The OFDM Paradox

To truly take advantage of OFDM, there needs to be a channel profile that can be changed over time to allow the system to learn and adapt to channel conditions.

Changing that profile requires good decision-making. The challenge is that the feature that potentially makes OFDM work better – the ability to optimize a profile – can also make it work worse due to the decision-making process coupled with operational realities.

For a single-profile downstream, when one CM has a channel problem, the impact of the solution to that problem (changing the profile) is directly felt by all other CMs. For a multi-profile system, when one CM has a problem, the solution (moving the CM to a different profile) does not directly impact the remaining CMs.

A single-profile system works by providing the worst service to all CMs.

A multi-profile system works by providing the best service to all CMs.

This paper will pursue this latter line of thinking and examine how to manage multiple profiles.

New Spectrum Opportunities

With advances in optical and RF equipment, it is now possible to extend the spectrum of the plant further in both the upstream and downstream. This is of particular interest when more spectrum is needed and where higher throughput is required.

In these new areas of spectrum, the HFC plant may have more variation or more micro-reflections, and other impairments where OFDM and LDPC may prove quite useful.

Below cutoff:

Below the downstream cutoff frequency (750 MHz, 1 GHz), the plant is generally well engineered. Even so, the amount of plant SNR differences within a downstream can be as much as 8 to 12 dB. This may allow for higher order modulation for CMs that are in areas of high SNR, such as homes after the first amplifier, compared to homes after the last amplifier where there is higher culmative noise.

Above cutoff:

To operate above the cutoff frequency, the plant will require upgraded amplifiers and optical nodes. This effectively moves the cutoff frequency higher. The only element that does not get upgraded is the taps. Most new taps have a viable frequency response up at least to 1 GHz and often past 1.2 GHz. Thus, OFDM would be very useful for a service from 1 GHz to 1.2 GHz where tap performance will cause channel variations.

QUANTIFYING THE VALUE

Channel Study

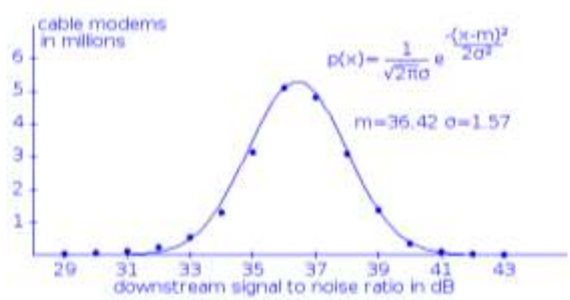

Figure 2: SNR Variation in an HFC Downstream

A study conducted by Comcast [1] and shown in Figure 2 measured the downstream

SNR for a large population of CMs on an HFC plant. As a general conclusion of this study, the consensus was that in an average HFC plant, there would be an 8 dB variation in SNR between various CMs.

The primary contribution to this variation has to do with the location of the CM along the HFC plant (topology constraints). CMs that connect to the coax segment nearest to the optical node and before the first amplifier see higher SNR values. A CM located at the end of the longest coax link after the last amplifier (typically 5 amps but could be more) typically has lower SNR.

With different profiles, different groups of CMs could be provided with a higher order modulation (where SNR is higher) while other CMs are provided with a lower order modulation (where SNR is lower).

A basic example use of profiles to study would be the following:

- Profile A: 256-QAM (2%)
- Profile B: 1024-QAM (25%)
- Profile C: 2048-QAM (64%)
- Profile D: 4096-QAM (9%)

In this model, the predominant modulation is shown. The model is also hierarchical where each level is progressively better than the previous level.

In parenthesis is the percentage of CMs from the study [1] whose maximum receive capability matched one of the profiles.

So, if a single profile were used, then all CMs would have to use the 256-QAM profile or the 1024-QAM profile (if the plant were upgraded to eliminate 256-QAM). With multiple profiles, over two-thirds of the CMs can have profiles that exceed the baseline profile.

Efficiency Analysis

OFDM Profile (percent of tones at each modulation level)

	QAM Levels					CM
Profile	4096	2048	1024	256	unused	Distribution
D	100%				0%	9.0%
C		100%			0%	64.0%
B			100%		0%	24.5%
A				100%	0%	0.0%
					non-D3.1:	2.5%

Throughput with a Spectrum of: 192 MHz

Profile	4096	2048	1024	256	Total Mbps
D	2,304				2,304
C		2,112			2,112
B			1,920		1,920
A				1,536	1,536

Aggregate Raw Throughput (Mbps):	2,029
Spread Highest to Lowest:	50%
Increase compared to worst-case channel:	32%

Figure 3: Throughput with Pure OFDM Profiles

OFDM Profile (percent of tones at each modulation level)

	QAM Levels					CM
Profile	4096	2048	1024	256	unused	Distribution
D	60%	25%	10%	5%	0%	9.0%
C		60%	30%	10%	0%	64.0%
B			80%	20%	0%	24.5%
A				100%	0%	0.0%
					non-D3.1:	2.5%

Throughput with a Spectrum of: 192 MHz

Profile	4096	2048	1024	256	Total Mbps
D	1,382	528	192	77	2,179
C		1,267	576	154	1,997
B			1,536	307	1,843
A				1,536	1,536

Aggregate Raw Throughput (Mbps):	1,926
Spread Highest to Lowest:	42%
Increase compared to worst-case channel:	25%

Figure 4: Throughput with a Blended OFDM Profiles

The data in Figure 3 and Figure 4 assume four profiles where each profile has a different percentage of bit loading. Figure 3 uses the rule where all subcarriers use the same modulation whereas Figure 4 modifies the modulation of each subcarrier to match the SNR of the plant at that frequency. Then a percentage of the total CM population is assigned to each profile. (Input data values in the study are not measured values)

Note that a higher average throughput is maintained even though a notable percentage of CMs are in lower performing plant segments.

As an example, a single profile implementation may have had to go down to 1,536 Mbps throughput whereas a multi-profile implementation was able to achieve 1,926 Mbps aggregate throughput. That is over a 25% improvement in raw throughput.

BUILDING THE MECHANISM

Codeword Builder

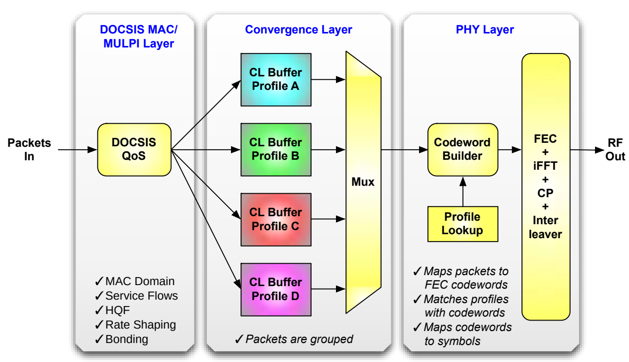

Figure 5: Codeword Builder

The OFDM channel generates a bit stream to the convergence layer. The convergence layer groups these bits into FEC codewords. DOCSIS encapsulated packets are then placed into the payload of these codewords.

Each path, as designated by a profile, is multiplexed at the codeword level. That is, each consecutive codeword can belong to a different profile.

An example implementation is shown in Figure 5. Packets exit the DOCSIS QoS process. As part of this process, they receive a tag that indicates what profile they are associated with.

The codeword builder collects these packets and sorts them into shallow collection buffers depending upon their tag. These buffers allow the PHY to accumulate enough bytes to fill a codeword. Packets are mapped directly to codewords and may be split across codeword boundaries.

A full-length codeword might have a payload of 1777 bytes and a total size of 2025 bytes. Thus, each codeword can contain on the order of one full length DOCSIS frame.

Latency

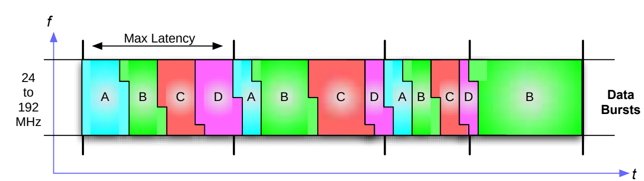

Figure 6: Profiles in the Downstream Path

The codeword builder can multiplex codewords in any manner it wants. This gives it the freedom to dynamically allocate data capacity to each profile depending upon the packet arrival rate.

The codeword builder does not actually care about data capacity which is pre-calculated by the rate-shaping mechanism in the DOCSIS MAC layer. Rather, the codeword builder is focused on minimizing latency.

Latency is dependant upon how often a path can be serviced. So, if a basic round-robin scheduler were used, as in A-B-C-D, then each channel would have to wait at least 4 codewords. This is shown in Figure 6.

If one path, say path B, has a higher bandwidth demand, the codeword builder may schedule A-B-C-B-D-B where the B path is serviced every second codeword.

Let's calculate how long it takes to send a codeword. Then we can figure out typical latency in terms of codewords. Assuming:

- 24 MHz min OFDM channel size
- 192 MHz max OFDM channel size
- 1024-QAM (10 bits/s/hertz)
- 16200 bit (2025 byte) codeword

The latency is calculated by:

latency = (codeword bits) / (bits/s/Hz) / BW

Then the time for one codeword varies from 8.4 µs minimum to 67.5 µs maximum. For four profiles at one codeword each, the latency per profile would be 34 µs and 270 µs. This latency would be in addition to any latency created by the downstream interleaver.

This leads to the following observations:

1. *There is a trade-off between the max number of profiles and latency.*
2. *This trade-off depends upon the scheduling algorithm of the codeword builder.*
3. *Larger channels provide lower latency and could accommodate more profiles for a given latency budget.*

SYSTEM OPERATION

Details

Profiles and OFDM Channels

Let's review the definition of a profile. First, an actual profile contains the dynamic configuration values for an OFDM channel. This configuration information includes modulation level for each subcarrier. A zero modulation level would imply a muted sub-carrier.

For convenience, each profile is assigned a letter. So, there will be Profile A, Profile B, etc.

Each downstream OFDM channel will have a range size, such as a minimum RF bandwidth of 24 MHz to a maximum of 192 MHz. Thus, the whole downstream 1 GHz spectrum could be covered with about five OFDM channels.

Each OFDM channel has its own unique set of profiles.

Thus, OFDM channel 1, Profile A is different than OFDM channel 2, Profile A. The reason for this is simple. Profiles describe subcarriers at a particular frequency, and each OFDM channel occupies different frequencies.

Profiles and Paths

From a packet forwarding perspective, each profile creates a unique path through an OFDM channel from the CMTS to the CM.

Packets travel paths. That path is described by a profile.

Thus, the two words "profile" and "path" often get used interchangeably when talking about how packets are sent from a CMTS to a CM across an HFC plant.

The term "profile" is a PHY level description; the term "path" is a MAC level description.

Multiple Paths

A CM may receive packets on more than one path. Thus, there are multiple forwarding paths, each with its own profile, from the CMTS to the CM.

Not all CMs have to receive on all profiles. One CM may receive profiles A and B while another CM may receive profiles A and C.

The CMTS will keep track of all the active paths to all the CMs. The CMTS will have to manage these multiple paths and determine through policy and forwarding rules what packets are to be placed on what path.

Path Assignment

The rules for profile usage are simple. The CMTS may use any path/profile for any task. The CM must accept anything addressed to it by any active path it may have.

For example, Profile A could be the common profile that all CMs can receive. It would be used for booting since the performance characteristics of new CMs are not known. Profile A could also be used for most or all MAC Management Messages (MMM), since it is by design the most robust and common profile.

However, the CM must be able to accept a MMM on any profile. This actually keeps things simple for the CM. The CM uses multiple paths and profiles to ensure packets get delivered through the PHY. It then uses MAC filtering to determine what to do with the packets.

This rule set also allows policies to change over time at the CMTS without impacting the implementation of the CMs in the field.

Service Flows

A service flow (SF) is a collection of packets that match a classifier. A typical classifier consists of the source and destination IP address, source and destination port, and the protocol type (this is the classic "tuple").

The most common SF is the default SF that contains anything not in a specific SF. Voice over IP usually has its own SF per voice call. Managed Video over IP will likely also have its own SF per video flow.

The convention followed will be that all SFs are fully contained within a profile for a given OFDM channel. That means that if a video flow is assigned to profile C in a given OFDM channel, then all packets in the flow remain on profile C.

Bonding

Since SFs are assigned to a single profile, there is no need to bond across profiles within an OFDM channel.

Bonding can be used to allow a SF to traverse multiple OFDM channels. For example, if a 5 Gbps SF was desired, that SF might be mapped across three OFDM channels.

Within each OFDM channel, a profile would be chosen. Since profiles are unique per OFDM channel, it could be any profile that the CMTS chooses. So, in this example, the three profiles across each of the three OFDM channels could be B-B-B or B-C-D.

Bandwidth Management

MAC or PHY?

So where does the decision get made as to what profile a packet should belong? Where does the bandwidth get calculated?

The MAC does all the bandwidth management. The MAC figures out if the packets will fit in the OFDM channel. The PHY packages up the packets and sends them out.

Thus, in theory, the packet buffers in Figure 5 should never overflow because the MAC will only send the packets to the PHY if there is room for them to hit the wire. That allows these buffers to be shallow, low cost, and low latency.

Channel Capacity

The aggregate capacity of the OFDM channel depends upon the configuration of each of the profiles contained in that OFDM channel and the relative usage of each profile.

There is no limit set on how much data capacity a profile may have. Any profile could be assigned 100% of the OFDM channel capacity at any time. The only rule is that all the profiles share the aggregate channel capacity. They can share that data capacity anyway that the CMTS sees fit.

If a profile is updated, then the channel capacity will change. This is a slow change since profiles are not updated frequently. However, when a profile is updated, the instant in time of the update to the rate-shaping mechanism must be synchronized to the instant in time of the update of the profile.

Time-Varying Channel Capacity

To calculate the throughput of an OFDM channel, you need to know the RF bandwidth and its modulation order or bit density.

This presents a challenge. Each profile has a different modulation order definition. Further, the amount of packets per profile depends entirely upon how many CMs are on each profile, the CMTS forwarding policy for packets, and the instantaneous traffic profile per CM.

That means that the bandwidth per profile is time-varying. It follows that the total channel bandwidth added across all profiles is also time-varying. In fact, if the modulation profiles ranged from:

- Profile A: 256-QAM (8 bits/s/Hz),
- Profile D: 4096-QAM (12 bits/s/Hz),

then the data capacity could vary as much as 50% at any point in time.

The MAC downstream rate-shaping mechanism must take this into account. To do this, the MAC downstream scheduler keeps track of which profile each packet has been assigned to and what the average bandwidth per profile is. It then calculates how much bandwidth is available on an instantaneous basis.

This is different than in DOCSIS 3.0. In all the early versions of DOCSIS, the max data capacity of the physical channel was constant. Now it is not.

Frequency-Varying Channel Capacity

The previous time varying calculation used the average throughput of a profile. But what if the modulation values of the subcarriers are different over frequency? For example, in the lower half of an OFDM channel, a profile may use 4096-QAM but in the upper half it may have dropped back to 1024-QAM.

From this, an average channel capacity can be calculated. But that average is only valid if the distribution of payload bits within a path is uniform across the entire channel. The channel is up to 5 codewords per symbol wide. If the profile definition is changed per codeword then it depends upon where that codeword lands within the channel and thus as to what the actual resulting data capacity of the profile is.

The difference between the predicted bandwidth by the rate-shaper based upon the average profile data capacity and the actual data capacity that results can be considered an error factor. This error can cause a slow drift in buffer levels. If there is less actual data capacity than calculated, then the PHY buffers may build up.

In practice, there should be enough randomness and buffering in the system that this averages out. A quick fix is to allow for this error in the rate-shaping calculations by subtracting out a few percentage points from the total data capacity. A good system design should also have a way of checking and correcting for this error.

Profile Management

How are profiles managed? When do profiles get updated? How do they get updated? Which CM goes to which profile?

For a profile to be useful, it should have the following measurable characteristics:

1. *Each profile should have a measurable and significant difference from another profile.*

2. *Each profile should serve a measurable and significant number of CMs.*

As the number of profiles increase, the system complexity tends to increase. Profiles are meant to be used sparingly and they should produce measurable and tangible results.

Static Profiles

The simplest solution would be to set all modulations levels within a profile, and perhaps even all amplitude levels, to the same value. For the small percentage of subcarriers that get into trouble, the LDPC feature of DOCSIS 3.1 will correct for a certain amount of errors.

For example, Profile A could be 256-QAM for all subcarriers. Profile B could be 1024-QAM for all subcarriers, and so on.

This approach is simple. It also creates and guarantees a strict hierarchy among profiles. Each profile will have a defined performance. If a CM does not work on one profile, it can be downgraded to a lower profile until it does work.

This is a good default operation mode and one that is predictable enough that it can be deployed with good success.

Dynamic Profiles

Despite the fact that static profiles will work, the flexibility of OFDM is almost completely lost. It seems like with a good set of test and measurement tools, these profiles could be updated, and performance could be optimized. Further, actual field problems that cause trouble tickets and truck rolls could be found and adjusted for automatically.

This is a tricky business. As mentioned before, it is easier to downgrade a profile than improve it. Thus, the algorithm should make sure that profiles do not degrade over time.

The algorithm also has to decide what to do when a CM is in trouble. Should it move the CM to a different profile, or should it update the profile?

One answer is that if the trouble is a single CM, then the CM gets moved. If the trouble is seen across a large number of CMs within a group, then the profile should get updated.

This leads to questions like:

- How much time should be used to detect common errors across CMs? The longer the measurement time, the more CMs with common errors that may be discovered. However, it also takes longer to fix the problem.

- How many CMs constitute a significant group to where the profile would get updated instead of moving CMs? Five? 10? 50?

- What is the nature of the problem? Is it frequency specific or is it channel specific? A clever algorithm must have the right information in order to make a judgment call.

If there is specific interference, such as interference from LTE, the best approach may be the nulling or lowering of modulation on specific subcarriers. If the problem is across the entire channel, then the SNR that a particular CM sees may have changed and the profile is no longer valid for that CM.

Measure and Sort

An ideal system with dynamic profiles would have the following capabilities:

1. The ability to measure the SNR for each subcarrier at each CM. This could be as many as 7680 measurements, one for each subcarrier for a 192 MHz OFDM channel with 25 kHz spacing. For CMs with multiple channels, there could be 2x to 5x the number of measurements (up to 38,400 values per downstream).

   This measurement would occur at boot time and then again at scheduled intervals or when there is a problem.

   Note that many of the measured SNR values will be the same. Thus, the resulting measurement data can be easily compressed prior to transmission.

2. The ability to predict from the SNR measurements what profile to assign a CM to.

3. The ability to test how well the CM performs on that profile.

4. The ability to analyze error scenarios for root cause.

5. The ability to fix those problems either by adjusting the profile or moving the CM.

6. The ability to do all this and yield a system that is efficient yet is also stable and reliable.

CHOOSING THE NUMBER OF PROFILES

Earlier, it was discussed that fewer profiles could provide a simpler system with lower latency. In this section, we will discuss use cases that impact the number of profiles required. We will then reconcile these two needs with a proposal.

The question is really separate for the CMTS and CM. A CM only needs to know about profiles that impact the CM, whereas the CMTS needs to know about all the profiles that impact all CMs. Thus, the CMTS probably has to support more profiles than a CM does.

HFC Plant Variations

MMM and Data

MAC management messages can be either multicast or unicast. Either way, it is important that they be delivered in a robust manner. The MMM traffic is also low bandwidth, allowing the MMM to be sent on a common profile like profile A.

Data is largely unicast and is thus sent directly to a single CM. It makes sense therefore to send data on the highest profile available to the CMTS and CM.

This requirement suggests that the CM must support at least two profiles.

Penalty Box

Sometimes there is a change in the HFC plant or even a home coax network that causes the SNR to be enough out of whack that the CM will either not work or partially work with lots of errors.

The right answer is a truck roll to fix the problem. However, prior to that truck roll, would it be possible, and of value, to offer an error-free lower speed of service to that subscriber until the problem can be fixed?

That is the concept of the penalty box. It would be a lower-order profile into which a CM would be sorted to when it cannot work on the higher order profiles. Once the CM is in the penalty box, the cable operator would be alerted so that the problem could be pro-actively fixed before the customer complained.

In the earlier example, the penalty box profile would be satisfied by the 256-QAM profile since the expectation is that the plant should work with 1024-QAM. However, in some cases, it might be necessary to add, say, an additional 64-QAM profile.

The penalty box may require support of an extra profile on the CMTS.

Modulation

In the generic example in this paper, four modulation levels were suggested. These were 256-QAM, (512-QAM skipped), 1024-QAM, 2048-QAM, and 4096-QAM. Each increase in modulation order (each extra bit/s/Hz) requires an extra 3 dB of SNR. Thus, this range of four modulation orders covers a 12 dB variation in plant. That is 50% more than the target 8 dB variation mentioned earlier and thus should be enough. Note that four distinct profiles could be defined with only a 9 dB variation in SNR if no modulation orders were skipped.

Inevitably, questions arise about also supporting 512-QAM. Additionally, there may also be future modulation orders such as 8192-QAM and 16384-QAM. These latter modulation orders may only work for short passive coax plants with home gateways.

Despite these additional modulation levels, it does not seem realistic that every modulation level would be needed within a particular fiber node's service area.

This requirement suggests that the CMTS must support at least four profiles and potentially a few more.

Geographical Differences

Since an HFC plant spans a geographical area, it is feasible that one part of a plant may be subject to interference that another part of the plant is not. This would have to do with the physical location of an interference source and how it couples into the HFC plant.

Let's say that a plant spans the east and west side of an area. Let's assume that there is an interference source on the east side of the plant that does not exist on the west side. Thus, the west side of the plant may have an average SNR that is, say, 6 dB higher than the east side.

Should there be separate profiles for the west side and east side? Well, it depends upon the fix.

If the fix is to not null out subcarriers and instead just run a lower modulation order, separate profiles are not needed. What may be needed is a large range of profiles at the CMTS.

If the fix is to null out a subset of subcarriers on frequencies where the interference is, then there may be benefit of using separate sets of profiles for the west side and east side.

The system would have to have an accurate way of measuring the interferers and of determining a set of suppressed carriers that works for all the CMs.

A simpler solution may be to use one set of profiles, but null out the subcarriers so that both the west and east sides are impacted equally.

This requirement suggests that if the CMTS could support more profiles, it could help with plant management. This would not impact the CM since the CM is on a specific plant segment.

IP Multicast

For IP multicast to work, all the CMs that are subscribed to a common multicast must be on the same path. That means a common profile.

The easiest solution would be to put multicast on the default profile. In this case, profile A would be used. By plan, all CMs can receive profile A, so multicast will always work.

This has the negative effect of multicast running slower than the unicast flows that exist on higher order modulation profiles. If IP multicast is not going to be deployed in any serious amount, then perhaps this is okay. At the same time, it becomes a self-fulfilling prophecy. If IP multicast is designed to always run slower than IP unicast, it is less likely to get used.

Since CMs are capable of receiving multiple profiles, a more clever algorithm could assign the IP multicast stream to the highest order profile that is currently shared among CMs. So if CM 1 supported profiles A-B-C-D and CM 2 supported profiles A-B-C, the multicast could be put on profile C. If a third CM that only supported profiles A-B joined, then the CMTS has the option of moving the multicast to profile B, or duplicating the multicast on profile B.

This requirement suggests that if the CM can support multiple profiles, the CMTS has complete flexibility in implementing multicast, and there is a lower or even no throughput penalty for using IP multicast.

Profile Update

How do you change a profile without risking a large disruption on the network if something goes wrong?

Whole Profile Update

The easiest way to update a profile is to issue a new profile and then hit the update bit on the profile descriptor to tell the CMs to start using the updated profile.

A more cautious way would be to change a small set of subcarriers at a time. The idea would be that the change would be minor enough that the CM FEC would still correct it, but the CMTS could discover that the FEC had to correct errors.

Migrate to a New Profile

Another approach would be to establish the new profile and then move CMs over to it over a period of time while measuring performance.

This requirement suggests that the CMTS and CM would need to support a spare profile.

Profile Changing

When a CM changes to a new profile, there is always a chance that it will miss some packets that got stuck in a CMTS queue. There are two basic solutions to this problem.

Pause / Play

Similar to the RF channel change techniques in DOCSIS 2.0, one approach for a profile change is for the CMTS to shut down the flow of packets to the CM, move the CM to a new profile, and then restart the flow of packets once the CM has properly synced to the new profile.

This design approach was taken because the DOCSIS 2.0 CM could only receive one downstream RF channel at a time.

This is an undesirable solution because CMTSs do not like to halt the flow of packets. Halting the packets of an unknown flow requires buffering. Buffering can overflow.

Add / Drop

A more elegant solution would be for the CM to first add the CM to the new profile so that it is receiving the old and new profiles. The CMTS would then move the traffic from the old profile to the new profile. Once this has occurred, and the buffers have emptied, then the CMTS can remove the CM from the old profile.

This technique may require sequencing of packets to ensure that packets are not delivered out-of-order.

This requirement would indicate that the CM supports one extra profile.

Probing

How can it be determined that a CM can be moved to a higher profile?

Predict

If the SNR of the CM is known, then it can be predicted what modulation level the CM can support.

The problem with this approach is that there are really 7680 SNR levels per OFDM channel. That is, the SNR varies with frequency. In addition, the LDPC FEC is very good at fixing problems where there are a small number of subcarriers that have SNR that is too low for the assigned modulation.

Thus, a good algorithm is required that can look at the average SNR and take into the account the strength of the FEC when making decisions.

Test & Measure

A more reliable technique would be to get the CM operational on a known profile, say, profile A. Then, in the background, tell the CM to receive packets on another profile, say, profile B, and to measure and report FEC and CRC error statistics.

The other profile could even carry a specific bit pattern that the CMTS generates so that the measured results are statistically accurate.

This requirement suggests that the CM needs to support at least one extra profile.

The 4+1 Proposal

Note that supporting more profiles does not mean that there are increased throughput requirements, since even one profile could be set to receive the full channel bandwidth. Rather, multiple profile support is about feature support.

It can be demonstrated that there are many reasons why it is useful for the CM to support more than two profiles. Many features require the support of at least three profiles at the CM. Simultaneous support of several features, and the desire to allow features to operate independently suggest an even higher number.

One proposal would be to support four active profiles and one spare profile. The spare profile could be used for temporary things like profile changes, probes, etc., while the other four would allow full, simple, and reasonable support of multicast.

The spare profile is like a spare tire. It is used only when changing something.

SUMMARY

OFDM allows for customization of a frequency profile to allow for optimum throughput. This profile could use the optimum modulation order for each subcarrier and would mute subcarriers where there was interference. LDPC is an error correction technology that effectively allows even higher orders of modulation to be used and allows for an “average” SNR to be used for setting modulation levels.

Despite the higher degree of flexibility that OFDM has, the HFC plant is fairly well characterized and maintained, and the amount of variation among CMs is minimized. Thus, a small number of profiles, such as four, may be sufficient to allow a compromise between system optimization and retaining simplicity.

REFERENCES

[1] Dave Urban, "The Case for Multiple MCS," **P8023.bn EPoC PHY Task Force**, 2012, http://www.ieee802.org/3/bn/public/adhoc_mmp/adhoc_mmp_urban_01.pdf

THE SHIFT FROM ELEMENT MONITORING TO PACKET MONITORING TO SERVICE MONITORING

Ian Wheelock, Hari Nair
ARRIS

Abstract

DOCSIS networks have always had good instrumentation for cable plant monitoring. The cable modems in the plant also function as probes for the network providing the operator good visibility into physical plant conditions. However as DOCSIS networks become the ubiquitous access network over which a wide variety of media and communications are delivered, and subscriber expectations of quality of service have been raised, operators have wanted to expand their understanding of network conditions from physical metrics to the quality of service delivered to customers.

In this paper, we describe a method to expand the role of DOCSIS Customer Premises Equipment (CPE) from physical and packet- level monitoring probes to aiding in service monitoring. We describe the considerations in the choice of active probes, and the network and service impact considerations from performing these tests.

EXISTING ELEMENT MONITORING

Current DOCSIS CPE provides a number of different measures that can be routinely sampled to develop an overall picture of how well the physical cable plant is operating. At a minimum, signal strength and downstream SNR values can be collected by the CPE, while upstream signal strength can be collected by the CMTS. Recording the actual downstream and upstream channels in use by the CPE allows for a network-wide view of physical cable plant performance to be accumulated over time. Maintaining time-based records of plant performance helps with identifying trends associated with plant health.

Metrics relating to received and transmitted packet and octet counts can also be sampled. Such sampling may be performed by a management system that collects details from either some or all of the CPEs and CMTS systems deployed in the cable plant. Combining the collected metrics with metadata associated with the CPE (such as traffic tier) and CMTS information (such as upstream and downstream channels) allows the management system to provide useful statistics as to the overall data health of the DOCSIS network.

Extensions like IPDR on the CMTS enable management systems get access to accurate byte counts for individual CPEs with granularity down to the DOCSIS service flow level. These byte counts can be used to determine overall usage levels of each individual subscriber, helping with limiting excessive usage or feeding into billing systems for per MB level billing.

The accumulated data metrics can also help with capacity planning. The constant 24-hour a day monitoring of the data performance of the network on an upstream and downstream basis (or on a more logical notion of DOCSIS 3.0 downstream or upstream service group (DS-SG or US-SG)) provides the actual data use versus the capacity of the network. Network management systems perform analysis of these figures and present information to the operator identifying problem areas of their network that are running above acceptable MSO performance thresholds. Such analysis allows MSOs to plan for capacity expansion of their networks in order to keep up with the data demand from their subscribers.

VoIP related monitoring is also provided by some management systems based on information available from deployed DOCSIS MTA devices and CMTS platforms. Such management systems help with the complete VoIP lifecycle management, from initial network analysis and optimization to ongoing fault, performance and capacity planning. Mean Opinion Score (MOS) data collected from MTAs is used in the analysis of the service.

Utilizing all these different types of monitoring provides MSOs reliable and useful information as to the performance of their network. The metrics are provided to the operators through suites of analysis and graphical tools. Some management systems tie these basic metrics together and generate summary ‘performance at a glance’ metrics, arranging the most likely problem areas in optimum ways for the MSO to act upon.

Even with all this data collection and analysis, the only real customer “service” that is being actively monitored is VOIP. Data bandwidth levels are monitored, but only really tell a small part of the customer experience. With so many services being delivered either by the operator (basic web serving, QAM based video service, walled garden IPTV service, WiFi service, etc.) or 3rd parties across the operator network (basic web serving, OTT video, gaming, storage backup, etc.), getting visibility into how well these services are performing from a subscriber viewpoint is critical to reduce churn and optimize the network.

ADVANCED NETWORK PROBES

Some systems do exist in the market to cover areas such as QAM based monitoring and IP service monitoring. These systems typically require a physical ‘probe’ unit to be deployed either in the operator’s network (in the case of end to end MPEG video monitoring) or by technicians into a subscriber’s home location.

MPEG video monitoring probes in the operator network are deployed and maintained in-situ as part of the overall video monitoring setup. They are in constant use providing realtime status information as to the quality of delivered video streams and associated metadata (SI, PSIP, etc). Multiple probes deployed in series provide operators with visibility in to how well the MPEG video streams are operating before they hit the final EQAM devices for transport over the HFC plant. Such monitoring enables faster troubleshooting in the event of a video outage, ruling the network in or out as the source of the outage.

Probes can also be used by technicians to troubleshoot issues occurring during installation, or alternatively can be deployed into a subscriber premises after problems post-installation have been reported. Providing the probe for remote troubleshooting is more cost effective than having a technician on site for hours or days attempting to observe a subscriber issue and then attempt to fix it.

New classes of probes are going beyond the basic analysis of MPEG video, and are now beginning to analyse the performance of IP, in particular IP video delivery. They are typically deployed as “man in the middle” probes, being able to observe all IP traffic of interest delivered to and coming from a subscriber network on a realtime basis. Monitoring traffic using this technique assists in determining the performance of such traffic, for example monitoring and reporting the bitrate associated with an adaptive bit rate (ABR) IP video stream being played by a subscriber device.

Deploying physical probes into subscriber premises attracts additional capex (probe cost) and opex (deployment) costs. These costs seem to be justified in terms of the results associated with using them (e.g. dealing with customer issues efficiently and keeping happy customers).

Now, imagine having the ability to enable equivalent probe functionality in the entire subscriber footprint of the MSO network without having to order a new piece of physical equipment or arrange a truck roll to deploy it?

MEASURING NETWORK PERFORMANCE

The FCC have been operated a program known as "Measuring Broadband America[i]" for the last number of years. The program was developed out of a recommendation by the National Broadband Plan (NBP) to improve the availability of information for consumers about their broadband service.

The goal of the program is to identify the nationwide performance of residential wireline broadband service in the United States. The program has concentrated on examining service offerings from the 13 largest wireline broadband providers using automated, direct measurements of broadband performance delivered to the homes of thousands of volunteers. Volunteers are provided with measurement devices known in the study as "Whiteboxes" that are configured to run custom software from SamKnows (a UK based company retained under contract by the FCC to assist in the program). The most recent report by the FCC contained measurements from over 6,700 Whiteboxes deployed across the footprint of the 13 selected broadband providers.

The Whitebox software runs a collection of periodic tests that aim to collect data representative of typical services in use by subscribers. Technical details of the 2013 report are available online[ii], but a summary is presented here.

The tests run by the Whitebox platform include download/upload speed, web browsing, UDP latency and packet loss, video streaming, voice over IP, DNS resolution and failures, ICMP latency and packet loss, latency under load, availability and consumption. The tests were conducted by the Whitebox devices to remote test nodes operating outside of the broadband providers networks. Test nodes were also deployed within the networks of the broadband providers and results were collected to those nodes, but were not included in the final results of the most recent (Feb 2013) report.

The February 2013 report, as in previous reports, emphasizes two metrics that are of particular relevance to consumers: ***speed*** and ***latency***. The report highlights the results of the following tests of broadband speed and latency

- Sustained download speed
- Sustained upload speed
- Burst download speed
- Burst upload speed
- UDP latency

The FCC program has become an important indicator of wireline broadband performance. A number of large cable MSOs actively participate in the study. The use of Whiteboxes in the network to conduct packet based testing and the acceptance of the results by the MSOs and wider broadband community indicates that active subscriber side measurement testing is extremely valuable and is considered representative of user experiences.

CREATING A DOCSIS CPE PROBE

Based on the brief technical description provided by the FCC on the operation of the

Whitebox device, it is clear that the actual measurements themselves are not overly complex. Deploying the same functionality (along with some of the DNS related tests) within existing DOCSIS CPE devices is well within the processing power of these devices.

Existing DOCSIS CPEs have considerable processing capabilities, allowing additional software to be developed to run in the home, such as fully featured routing, voice over IP call handling, configuration web page, etc. Some CPEs already have some basic troubleshooting options for assisting in technician installs, as well as for running local tests remotely from the NOC.

New and future DOCSIS CPEs are being developed using extremely capable SOC (system on chip) processors. The SOCs are incorporating large numbers of functions that formerly resided in external integrated circuits on the CPE motherboard (such as Ethernet switches, MoCA interfaces, voice processing functions, etc.). In addition to the extra functionality, the SOCs are getting much more processing power than before.

This processing power is being targeted at new functions like 'gateways' that provide a one-stop shop for all home networking needs, or as 'application servers' that provide new software services for the home (that would otherwise run on a dedicated PC in the home).

Some of this same processing capability can be utilized to provide monitoring functions in the CPE similar to those available in custom probes. Being able to download a firmware update with probe-like functionality to deployed CPE devices enables the entire network to participate in packet and service monitoring.

PACKET MONITORING

Based on the earlier descriptions of network probes and the FCC Measuring Broadband America program, the industry seems to have accepted the use of probes for troubleshooting hard-to-diagnose network issues, and for determining subscriber related latency and speed measurements.

Using modified DOCSIS CPE probes to perform these FCC-style tests provides the option to have visibility of performance across the entire DOCSIS network footprint. Running these tests and reporting realtime results can help identify real issues occurring before a customer call center is inundated with support calls.

Control Platform

Specialised control platforms are required to manage such large numbers of available probes in the network. Coordinating subsets of probes to run tests across the network requires advance knowledge of probe presence within the topology. In addition to running the tests, the actual results of the tests must be collected and stored. The stored results are then analysed to provide refined performance results to the operator in realtime. Analysis rules must be defined to limit the total information supplied to what can be used.

Overall performance data from every node in the network can be modeled. In the case of parts of the network reporting poor performance the control platform must be able to integrate with other management systems that are already performing element monitoring (for plant monitoring and data counts). The combination of the packet monitoring and existing element monitoring provides a powerful tool for troubleshooting issues in the network. Not only can the probes be used for periodic performance monitoring of the entire plant, they can also be brought to bear on issues identified by other management systems. Running actual traffic tests through problem areas of the plant can help diagnose issues.

The control platform needs to understand the location of test nodes that the probes rely on to conduct their tests. The FCC report relied on results from Whitebox to off-net test nodes. In the case of monitoring the MSO DOCSIS network, on-net test nodes can be used by the probes as well as off-net test nodes.

As the DOCSIS CPE probe is a software modified cable modem, activating this software has the potential of enabling probes at every point of the network. This is an extremely powerful benefit, but one that has to be manged carefully. Running performance tests on every modem in a network segment at the same time will completely saturate the data network. Careful probe selection across the network topology and scheduling of test runs, performed by the control plane platform, is required to minimize the overhead of traffic testing to the overall network. At a minimum selecting single probes on every DS-SG/US-SG would likely give sufficient coverage across the network. The speed tests could be organized to run periodically, say once an hour, but those tests can generate large bursts of traffic. In an effort to get more visibility per hour of collected results, more probes maybe required per DS-SG/US-SG. Some other tests, such as UDP latency, can run for much longer periods but only generate minute amounts of data. Scheduling tests to run across every hour while minimizing the impact on the network is crucial to getting realistic performance monitoring metrics.

Internal Network Monitoring

In most cases the on-net test nodes used by probes will be located relatively close to the CMTS the probes are associated with. However, it is also possible to consider using on-net test nodes throughout the MSO data network to provide additional network visibility. Having probes communicate across internal boundaries in the MSO network allows traffic routing performance to be monitored. Tracking the results of probe tests to these remote on-net test nodes can provide an indicator as to how stable the routing of the network is. Marked changes in latency or speed performance across these results can point to network equipment outages, incorrect route updates, DNS issues, or other subtle networking issues.

SERVICE MONITORING

Deploying DOCSIS CPE probes and using the limited set of FCC tests outlined earlier (along with DNS related tests) bring a new dimension to determining the performance of the network. Like existing element monitoring, results of the tests can also be retained and used for trend analysis to understand if changes to the network are resulting in realistic performance benefits for users.

The limited tests are concentrated in specific areas, though. They are focused on speed and latency and do not directly consider the performance of services, such as IP video or VoIP. The FCC-deployed Whitebox does include the ability to run additional tests more focused on these areas.

With the introduction of more powerful SOCs as outlined earlier, the newly released latest generation DOCSIS CPE platforms and gateways are much more capable of running software to mimic services such as IP Video streaming or testing VoIP services. As well as acting as service clients, these devices can also be modified to perform “man in the middle” analyses, monitoring service specific flows directed to connected home devices. Monitoring the actual TCP connections allows for bandwidth, delay, packet loss and other performance indicators to be identified in the home in a similar fashion to how some existing network probes operate.

These new platforms also provide WiFi and MoCA networking services to the home. Using the same “man in the middle” monitoring, it is possible to distinguish between the different network interfaces in use and provide targeted network measurements. Such measurements can be very useful, for example, in helping to determine degraded home WiFi performance.

Some of these platforms are hybrid devices, providing both MPEG QAM video tuners alongside DOCSIS tuners. Such platforms can be adapted to provide the typical QAM MPEG video analysis provided by standalone probes available today. Running the full set of ETSI TR101290 fault detection is doable within the processing power of the new SOC platforms. Spectrum analysis features within the DOCSIS platforms can also be combined with the fault detection to provide a complete QAM monitoring solution. As devices such as these are deployed within the DOCSIS network, they can start to act as “canaries”, performing QAM/MPEG video analysis on behalf of existing STBs in the network that are only capable of playing back video.

SUMMARY

The existing network management platforms in use today rely heavily on element monitoring of CPE and CMTS devices. The information provided helps with capacity planning and troubleshooting network problems. The broadband industry seems to accept the use of distributed probes to collect detailed performance data for their networks. Extending the DOCSIS CPE device with basic network test functionality effectively makes it a probe. Coupling these probe devices with new backoffice software to coordinate the probes across the DOCSIS network brings visibility to how the overall network is functioning. Comparing all parts of the network together helps identify where the network is running at peak performance, or where problems exist. Integrating the probe results/analysis with existing element monitoring helps to develop a bigger picture of how the network is running. Extending the probe capability in newer SOC platforms brings even more visibility into service monitoring. New hybrid devices bring the capability of QAM/MPEG video analysis enabling real end-to-end video monitoring wherever such platforms are deployed in the network. Being able to collect performance results of how well services are being delivered over different home network interfaces helps determine subscriber quality of service and quality of experience.

---

[i] http://www.fcc.gov/measuring-broadband-america

[ii] http://data.fcc.gov/download/measuring-broadband-america/2013/Technical-Appendix-feb-2013.pdf

WHAT'S HOT: LINEAR POPULARITY PREDICTION FROM TV AND SOCIAL USAGE DATA

Jan Neumann, Xiaodong Yu, and Mohamad Ali Torkamani
Comcast Labs

Abstract

Large numbers of TV channels are available to TV consumers these days. Such choice is both a blessing and a curse, because too many options can overwhelm the consumer and due to the limited screen real estate on devices, only a small number of programs can be presented at a given time. To address this issue, at Comcast Labs we work on algorithms that compute rankings of current and upcoming programs based on various relevance criteria.

In this paper we describe one of our algorithms, where we predict the future popularity of programs by combining information from historical Nielsen ratings, DVR scheduling activity, and social web activity (e.g. Facebook, Twitter).

INTRODUCTION

The question of "What's On TV?" is part of the daily ritual of watching TV. Usually we start by examining "the grid" and surf from channel to channel to find out what programs are playing on what channel. The order of the channels in the channel lineup rarely changes and though it is based on thematic groupings, it does not reflect that the themes and popularity of different channels and programs changes over the course of a day.

To address this issue, at Comcast Labs we developed an algorithm that predicts the popularity of programs that are currently on TV or will be playing in the next 24 to 72 hours. The output of this algorithm is then used to present schedule information to customers in order of (predicted) popularity of a given program and aims to give them an improved user experience (see Figures 1 and 5 for screen shots of Comcast Interactive Media's "What's On" iPhone app).

Figure 1 – Screenshot of "What's On" iPhone App

Currently, the most prominent metric to measure the popularity of TV programs and channels is provided by Nielsen Media. They publish the well-known suite of Nielsen TV ratings. One of the ratings, for example, is the percentage of TV consumers that are currently tuned to the program of interest.

During the last couple of years the consumption patterns of TV consumers are undergoing a rapid change where content is consumed on a range of devices such as cell phones, computers and tablets in addition to the TV. Also audiences nowadays tend to interact socially with TV programs via Twitter, Facebook and other social web sites and such activity can be utilized to further gauge the engagement of the audience with a program, as we will do in this paper.

On these sites viewers of a program indicate their level of like (or dislike) for it by publishing messages related to the program content or actors (e.g. Twitter), give explicit feedback via like/dislike buttons (e.g. Facebook), or even indicate that they are currently watching a TV program (e.g. Zeebox, GetGlue, IntoNow, Shazam, etc.).

Our approach uses machine learning to build a model that combines statistics about past Nielsen ratings, and scheduled DVR recordings, together with current social signal activity to accurately predict the popularity of one program relative to another.

Each of these sources of information captures a different notion of popularity. In the following sections we will describe the sources of information that we are using and the algorithm in more details

NIELSEN TV RATINGS

Nielsen ratings have been used by content providers for a long time to measure the audience participation of TV programs. The Nielsen shares, the percentage of viewers that are tuned to a given channel or program compared to all consumers that use their TV at the moment, are used to judge the success of a program and to set the rates for advertisers. Due to the nature of the data collection, Nielsen audience measurements are only available with some delay for most channels and programs since the viewing numbers also include the delayed consumption of programs on the DVR.

Nielsen national channel ratings are determined by monitoring the TV consumption behavior of a small sample of households and then extrapolating these sample statistics to the universe of all TV consumers in the US. We looked at the average number of TV viewers tuned to a given channel across the US for any 15 minute interval of a given day, which we will refer to as *Nielsen channel rating* from now one, as well as the average number of TV viewers tuned to a given program, which we will denote as *Nielsen program rating.*

We need to do some preprocessing steps before Nielsen ratings can be used together with the other usage data. Since a Nielsen channel or viewing source corresponds to a number of physical related channels (e.g. all NBC broadcast channels are aggregated in a single “NBC Nielsen channel”, all HBO-East, HBO-West and SD//HD channels map to a single “Nielsen HBO channel”), we semi-automatically created a mapping between the physical stations to the Nielsen aggregate channels.

Also, since Nielsen does not use unique ids to identify programs, we need to match a given program in the schedule to the corresponding program in the Nielsen ratings report. This is implemented using a combination of editorially created regular expression matching together with natural language based distance metrics. After establishing correspondence we find the ratings for the program at the same time and weekday for a fixed number of preceding weeks.

The same process is repeated for the channel popularity. In our experiments we utilized six months of Nielsen national channel and program ratings.

DVR SCHEDULED RECORDINGS

Using Nielsen TV ratings to predict popularity of programs that were never or not recently aired on TV is challenging, because there are simply not enough samples available for an accurate prediction. Examples of such programs are yearly awards shows such as the Oscar’s, Emmy’s, Grammy’s, large sporting events like the Olympics, NFL or NBA playoffs, news breaks, and newly scheduled programs

To be able to deal with such programs, we utilized the DVR scheduling statistics to count how many customers have scheduled their DVR to record a given program.

We compute the DVR score by aggregating the number of scheduled recordings for specific episodes as well series across all users that are stored in Comcast's online DVR scheduling service.

While doing so we make sure to account for differences in the number of customers in different markets, so that we arrive at a normalized DVR score that can be integrated with the remaining scores.

SOCIAL WEB ACTIVITY

There are two main types of information that can be gathered from social networking web services. One examines the connections between different participants in a social network (e.g. friends, followers), and the other looks at the activity between the participants in such networks. In the approach described in this paper, we only considered social activity based measurements since we are interested in aggregate popularity estimates, not personalized recommendations.

According to a recent Nielsen/SocialGuide study there is a strong correlation between the Twitter activity related to a program, as measured by tweets containing the hash tags associated with it, and TV ratings [1]. The study found that for young adults (14-34 years old), a 8.5% increase in Twitter activity correlated with a 1% increase in TV ratings for premiere episodes, and a 4.2% increase in Twitter activity correlated with a 1% increase in ratings for mid-season episodes. For older TV consumers this effect was weaker, but still present (i.e. a 3.5% increase in Twitter activity correlated with a 1% increase in ratings).

In contrast to watching TV, information about participation in social web services is usually made available to third-parties via APIs. For example, when a someone tweets a message related to a TV program, Twitter makes this message instantly available on its message feed and a third party can easily analyze and filter the information and make aggregate information available in real-time.

We used an external company to provide us with the aggregate counts of Twitter and Facebook activity for the time period a program aired on TV +/- 3 hours for various markets. See Figure 2 for some example data.

As with the DVR score it is important that the social activity signal is normalized with respect to the number of participants in

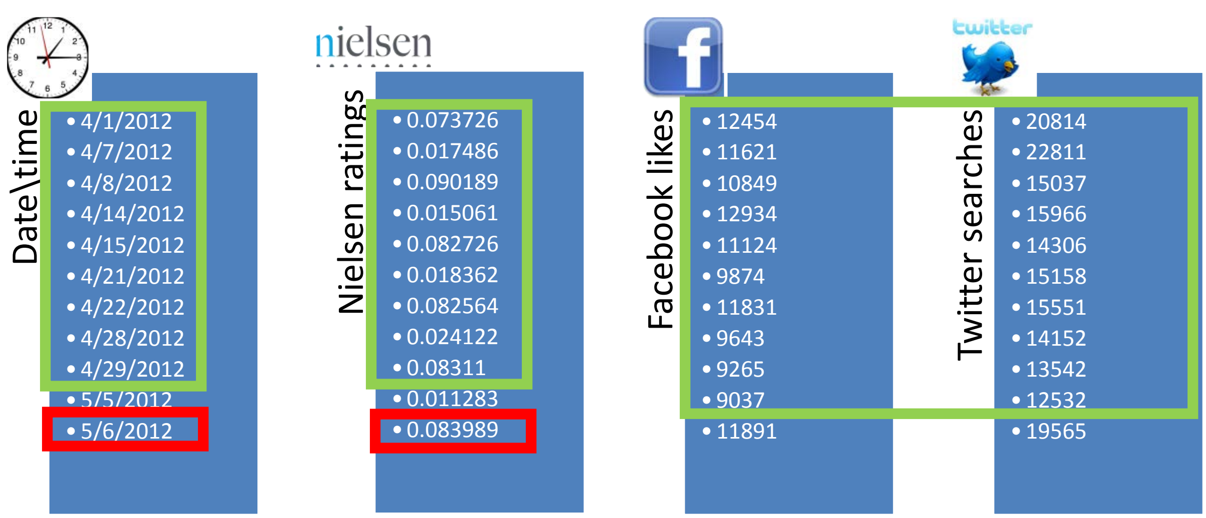

Figure 2 – TV rating prediction model: we use past Nielsen ratings, and current + past social activity signals associated with a program to predict its future rating (red is target value, green are input feature values for the regression model)

For each program and 30min time interval/date do the following:

1. **Extract sufficient statistics**: for each popularity score (Nielsen, DVR, Facebook likes, Twitter activity,…)
    a. Find Nielsen program scores for 7 last airings of a program, if program scores for less than 7 prior programs can be found, use channel scores instead.
    b. Find social signal and DVR scores for current program
    c. Compute the following statistics: Max, Mean, Median, Last value, Mean of the last 3 values, median of the last 3 values
2. **Model Estimation (only during training phase)**: Train a regression or classification function for past airings of this program for which we have data, use historic Nielsen program and channel scores as target variables.
3. **Prediction**: Based on the trained model predict the current program popularity.
4. **Ranking**: Based on the predicted scores, sort the programs.

Figure 3 – "What's hot" prediction algorithm

different geographical regions, so that the scores we use can correctly be used to predict popularity for a target distribution whose statistics differ.

TECHNICAL APPROACH

To predict the future popularity we have to build a model of how the different sources of information about customer activities predict future popularity. We start with the schedule for the upcoming 72 hours, identify all the programs for each station that are playing during each 30 min interval and collect relevant historical information for the different sources, e.g. Nielsen channel and program ratings, number of scheduled recordings of a given program, and the associated social activity signal.

Combining these different scores into a consistent ranking function is not straight forward, since not every score is available for each program, and scores differ in how much they change over time or correspond to different embodiments of user behavior. For example, the coverage of program ratings by Nielsen is only about a third of the programs that are scheduled for a given 24 hour period, while Nielsen national channel ratings are available for about 120 channels that cover 90% of the programs that are typically being watched. On the other hand, the distribution of DVR scheduled recordings is much more peaked, than the distribution of Nielsen ratings across programs. This is likely due to the fact that a customer only schedules a handful of programs for recording, while not being as selective while browsing the TV.

Using future Nielsen program and channel ratings as the target variable, we compute a range of statistics on each input, which is then used as a feature in a regression or classification framework to approximate the target variable as closely as possible.

This prediction component is then input into a temporal filtering framework to compute the final ranking function that is used to sort the programs. The full high-level algorithm that we implemented is described in Figure 3.

PREDICTION MODEL

We will start by defining the notion of a rank function. Our goal is to learn a function *f*, so that $f(x) > f(y)$ if program x is supposed to be ranked higher than program y. We explored a number of approaches to learn such a ranking function. The function *f* can be

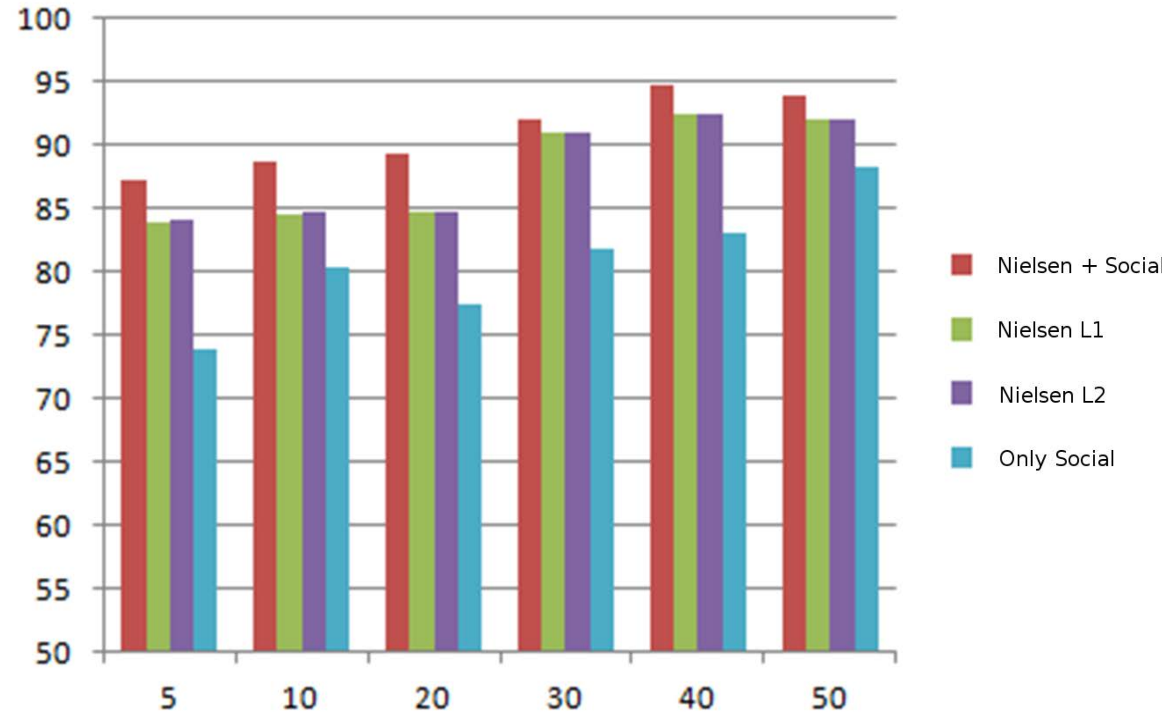

Figure 4 – Top k accuracy of "What's hot" prediction for different input sources

optimized in many different ways. We studied modeling the ranking problem as a pairwise classification problem, i.e. find a classification function that returns a positive value if *x* should be ranked higher than *y*, and negative otherwise. We also looked at regression functions to model the ranking *f* directly. For the classification approach, we explored support vector machines, k-nearest neighbor approaches, as well as a random forest classifier [2]. For the regression models, we looked at linear models (both with L1 (absolute value) and L2 (least squares) regularization terms, k-nearest neighbor regression, support vector regression, decision trees, and random forest regression [2]. At the end we got the best results using shallow random forest regression trees with past Nielsen scores and current and past social signals as feature inputs as described in the (see Figure 2 for an illustration of the input features and target variables).

EVALUATION AND RESULTS

To evaluate our system, we take the true Nielsen scores as our gold standard and evaluate our predicted popularity ranking against it. We used viewing data from June 2012 to train our predictors and predicted the ranking of programs for every 30 minute interval for the first week of July 2012. The evaluation criteria we are using is the Top-k criterion, i.e. how many of the top k programs of the ground truth data can be found among the top-k programs of the predicted data set. In our experiments we varied k from 5 to 50 in increments of 5 (see Figure 4). As described before, we got the best results using random forest regression trees, but we varied the set of input features that we considered. The results are summarized in Fig 4. One can see that social network activity by itself does not perform very well compared to using a moving least squares (L2) or robust estimation (L1) using a window of Nielsen ratings. If we combine both historical Nielsen ratings using our non-linear temporal filtering framework with random forest regression trees and all the social signals leading up to the show, then we can get a 4% increase in Top-10 accuracy over only using Nielsen ratings for prediction.

APPLICATIONS

The prediction of the most popular program for a customer has many applications. To name just one, some example screen shots of the "What's On" app, developed by Comcast Interactive Media, can be seen in Figures 1 and 5. This app allows a customer to see what is currently or soon showing on TV, sorted by different criteria such as most popular, favorite channels, movies, etc. The output of the algorithm can also be used in any other set-top box and mobile application where we want to return a popularity-ranked list to the customer.

CONCLUSION AND FURTHER WORK

In this paper we presented an approach that combined Nielsen ratings, DVR schedule information, and social networking activity measurements in a temporal filtering framework to predict the popularity of future programs. The experimental results showed that combining TV ratings with measures of social network engagement leads to more accurate predictions for relative popularity rankings of TV programs than just using TV viewership numbers alone.

The framework we described in this paper can be extended in a number of ways. Fore example, one could design more complex models to predict a program's popularity that incorporate both program related attributes and other non-TV measures of popularity. Examples of program attributes are indicators if the program is a new program or if the episode of interest is the season premiere, what genres a program is associated with, the actors in it, directors for movies, etc. Other

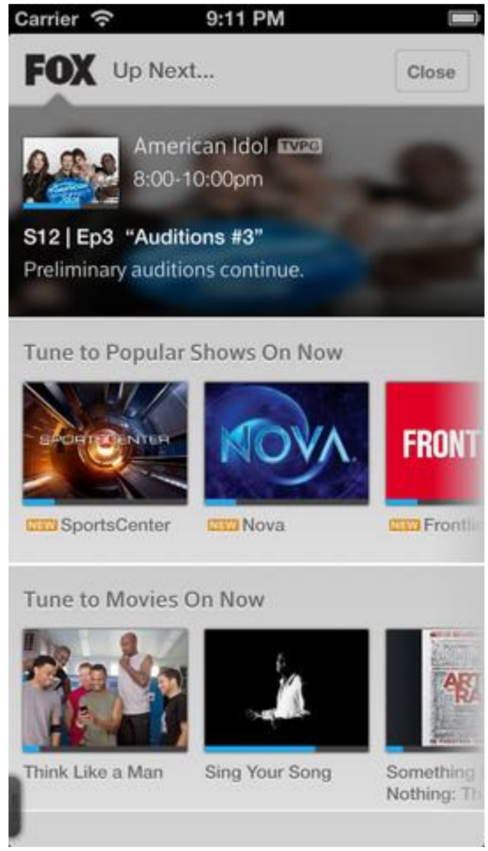

Figure 5 – Sample Client App Screen

measures of popularity we are looking at are box office numbers for movies, Rotten Tomatoes reviews [3] and even the presence or absence of editorial recommendations.

Finally, we are also looking at combining the aggregate popularity prediction described in this paper, with personalized recommendation algorithms that take a user's TV consumption history into account to deliver truly personalized TV recommendations to customers.

REFERENCES

[1] http://www.nielsen.com/us/en/newswire/2013/new-study-confirms-correlation-between-twitter-and-tv-ratings.html

[2] T. Hastie, R. Tibishirani & J. Friedman (2001), The Elements of Statistical Learning, Springer, New York, NY

[3] www.rottentomatoes.com

ISBN 0-940272-01-6; 0-940272-08-3; 0-940272-10-5; 0-940272-11-3; 0-940272-12-1; 0-940272-14-8; 0-940272-15-6; 0-940272-16-4; 0-940272-18-0; 0-940272-19-9; 0-940272-20-2; 0-940272-21-0; 0-940272-22-9; 0-940272-23-7; 0-940272-24-5; 0-940272-25-3; 0-940272-26-1; 0-940272-27-X; 0-940272-28-8; 0-940272-29-6; 0-940272-32-6; 0-940272-33-4; 0-940272-34-2; 0-940272-35-0; 0-940272-36-9; 0-940272-37-7; 0-940272-38-5; 0-940272-39-3; 0-940272-40-7; 0-940272-41-5; 0-940272-42-3; 0-940272-43-1; 0-940272-44-X; 0-940272-45-8; 0-940272-46-6; 0-940272-47-4; 0-940272-48-2; 0-940272-49-0; 0-940272-50-4; 0-940272-51-2; 0-940272-52-0; 0-940272-53-9; 0-940272-54-7