METHODS TO INCREASE BANDWIDTH UTILIZATION IN DOCSIS 2.0 SYSTEMS

Daniel Howard, Laura Hall, Keith Brawner, and Hans Hsu
Broadcom Corporation
Nick Hamilton-Piercy, Reynold Ramroop, and Sheng Liu
Rogers Cablesystems

Abstract

Several methods to increase bandwidth utilization in DOCSIS 2.0 systems are presented, and results of a field trial that demonstrated several of the methods are reported. On the downstream where the majority of packets transported are large packets, the use of wider bandwidth RF channels of 12 or more MHz is discussed, where statistical multiplexing gains improve the channel capacity above and beyond the factor by which the bandwidth is increased. Similarly, going to higher order modulation on the downstream (1024 QAM) also increases the downstream capacity, and in particular can provide MSO's the ability to get more HDTV channels per 6 MHz of downstream bandwidth or 25% more data capacity in DOCSIS downstreams.

On the upstream where small data and voice packets can be the majority of packets transported, bandwidth saving techniques such as dynamic header suppression and synchronous operation reduces the packet overhead that can account for a significant fraction of the minislots required to transport the packet. For medium and large packets on the upstream, the use of higher order modulation up to 256 QAM can be used to expand the capacity, and when combined with the ingress and impulse robustness features of DOCSIS 2.0 systems, 256 QAM upstreams can be supported on today's cable plants, even in bands which previously could not support more than QPSK. These facts are born out by the results of a field trial in which both 1024 QAM on the downstream and 256 QAM on the upstream were demonstrated, the latter in the presence of ingress noise on the cable upstream. The conclusion is that many cable plants are already capable of supporting higher orders of QAM modulation, and thus capacity increases of up to 33% above that already provided by DOCSIS 2.0 are possible.

INTRODUCTION

It is generally recognized that bandwidth needs of users on cable plants will continue to increase as new applications and higher quality versions of existing applications emerge. Applications that increase bandwidth needs on the downstream of cable plants include high definition TV (HDTV), video-on-demand (VOD), voice over IP (VoIP), and higher speed data service. The growth in data rate requirements over time is still considered exponential, doubling every year at least, however the trend can perhaps be better appreciated via applications that are known to be increasing the current data rate requirements for cable operators in particular. For example, activities such as more frequent and larger file transfers via email and web browsing due to imbedded high quality photographs, audio, and video will drive both upstream and downstream data rate requirements upward in the near future. In particular, the transfer of video files over broadband Internet connections is already part of current product lines for personal video recorder devices such as RePlay TV's current offering [1], where the capability to share movies over the Internet with friends and family is touted.

On the upstream, home based servers and peer to peer networking applications a la Napster (and its current look-alikes) will continue to drive bandwidth needs upward. And upstream bandwidth must be provided in a more robust manner, since RF interference is frequently present. When RF spectrum below 25 MHz must be used for additional upstream data channels, the modulation technology must be robust to ingress, impulse, and higher thermal noise conditions.

The upstream capacity has recently been addressed by DOCSIS 2.0 technology, which provides significant increases in upstream capacity and robustness on cable plants. Compared to the DOCSIS 1.1 maximum rate of 16 QAM operation at 2.56 Megasymbols/sec, DOCSIS 2.0 technology enables at least 64 QAM at 5.12 Megasymbols/sec, an improvement of three times in raw capacity. DOCSIS 2.0 provides this improvement by providing a new modulation technique, SCDMA, which is inherently robust to impulse noise, by increasing the robustness of TDMA via byte interleaving and increased FEC, and by the proprietary schemes most vendors provide for cancellation of ingress.

DOCSIS 2.0 does not address capacity increases in the downstream, however, and if data rate requirements are doubling every year, then in two years the three-fold increase in upstream capacity provided by DOCSIS will be used up and additional improvements in bandwidth may be needed.

Hence, there is still a need to consider techniques that increase capacity in both the upstream and the downstream on cable plants. In this paper, several techniques for increasing the capacity on cable plants are presented, and field tests of some of these techniques are also presented as proof of their viability. The techniques include higher order modulation on both the upstream and downstream, synchronous operation on the upstream, and dynamic payload header suppression on the upstream.

METHODS AND BENEFITS OF INCREASING CAPACITY ON THE DOWNSTREAM

Wider Channels

On the downstream of cable plants, there are two main techniques that can be used to increase the raw capacity of data services: increasing the total channel width, and increasing the order of modulation. It should also be noted that increased video compression via MPEG4 is another way to get more channels in the same RF bandwidth, however in this paper we focus on the media access control (MAC) and physical (PHY) layers.

Increasing the downstream channel width can be done in two manners: first, the symbol rate can be increased. Since cable downstreams are channelized on 6 MHz spacing in North America and 8 MHz spacing in Europe, the most efficient manner to use for increasing the symbol rate of DOCSIS downstream signaling would be in integer multiples of 6 MHz (or 8 MHz in Europe). Hence, the first method of increasing the downstream channel capacity is to increase the downstream symbol rate from about 5 Megabaud to 10 Megabaud and the subsequent downstream RF bandwidth to 12 MHz, or twice the current 6 MHz RF bandwidth.

Note that doubling the channel width does not necessarily increase the spectral efficiency of the transmissions since the alpha factor used in symbol shaping may remain constant. Even though the guard band in between the two individual channels is removed, more guard band on the edges of the signal spectrum is required in terms of Hz for the same value of alpha when a larger symbol rate is used. Since doubling the channel width does not increase the spectral efficiency, the main benefit of this technique lies in the additional statistical multiplexing gain that comes from wider channel widths. Essentially, leftover capacity in each of the individual channels from gaps in scheduling transmissions can be combined and used for additional transmissions in the single, wider channel. Further, latency can be reduced by exploiting earlier opportunities to transmit in the combined channel instead of waiting for opportunities in a single, smaller channel. Estimates of statistical multiplexing gain vary from 10% to 40% [2], depending on the traffic, size of the original channel, and the number of channels being multiplexed.

One of the issues associated with using larger channel widths on the downstream is that legacy modems cannot use the larger channels, being limited to conventional 6 MHz channels. Thus, the second method of increasing the statistical multiplexing gain on DOCSIS downstreams is to combine multiple downstream channels logically so that a modem can receive on multiple 6 MHz channels simultaneously. This method allows future modems to access larger channels, and headend schedulers to use leftover capacity in the individual channels more effectively, but at the same time permits legacy modems to continue to use the individual 6 MHz channels. A version of this technique was described in a recent NCTA/SCTE paper by ATT [3], where 40% gains in channel utilization from statistical multiplexing were shown for combinations of four downstream channels.

Higher Order Modulation

Wider downstream channels will clearly provide some level of increased performance on the downstream, however traffic variation and the ratio of new to legacy modems cause statistical multiplexing gains to be variable and difficult to predict for multiple and/or wider downstream channels. Another technique, which gives clear and predictable gain in the downstream channel, is the use of higher order modulation, for example 512 QAM or even 1024 QAM. In the latter case, the spectral efficiency is increased from 8 bits/symbol (256 QAM) to 10 bit/symbol, an increase of 25%. Further, by increasing the channel capacity directly, there will also result a statistical multiplexing gain in that channel. Figure 1 depicts the received constellation of a 1024 QAM downstream signal.

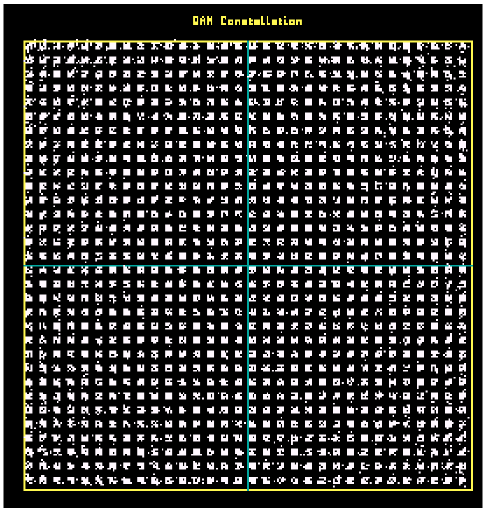

Figure 1. 1024 QAM DS Signal

But higher order modulation on the downstream has been criticized in the past based on the difficulty in reliably operating even 256 QAM on the downstream [4]. In particular, the cresting of CTB and CSO on the downstream and higher SNR requirements are often quoted as factors that could limit or prevent the successful operation of higher order modulation on the downstream.

For the SNR issues (which also arise in using double- and quadruple-wide downstream channels), it may be noted that in current cable downstreams, the transmit power of digital channels is backed off from the level used by analog channels, by up to 10 dB. The reason is that analog TV requires an SNR of up to 46 dB for high quality video, while a practical digital TV receiver requires only about 30 dB of SNR. In going to higher order modulation such as 1024 QAM, another 6 dB of SNR would likely be required, which is still 10 dB below analog levels. But if the noise floor of the plant was insufficient, clearly when all analog channels are replaced by digital carriers, the resulting digital carriers can be transmitted at higher levels due to the laser power margin freed up by removing the analog channels. And if only a subset of analog channels were boosted in order to support higher order modulation, the effect on overall plant balancing would be minimal.

The CTB and CSO issues can be dealt with using many of the techniques described in [4], examples of which include additional interleaving and/or coding, better equalization, and also offsetting higher order QAM frequencies to avoid the strongest CTB and CSO ‘tones’.

The fact that modern transmitter and receiver technology has mitigated many of these issues is born out by a field test of higher order QAM on the downstream, described in a subsequent section below. Further tests are planned.

Benefits of Downstream Improvements

While 25% improvement from higher order QAM and 10-40% from statistical multiplexing may not sound like drastic improvements in downstream capacity, consider the benefits for an application that is currently in the headlines for cable operators: HDTV. Currently in a 6 MHz downstream channel using 256 QAM, cable operators can deliver 2 High Definition (HD) channels without degradation, 3 HD channels using statistical multiplexing and allowing a slight degradation in quality. If the channel width were doubled to 12 MHz, then a 40% statistical multiplexing gain would mean that 5 channels could be delivered in 12 MHz with no degradation, and a 24 MHz wide downstream channel could deliver 11 HD channels.

But now consider increasing the order of QAM to 1024 on the downstream. The additional 25% raw capacity plus an additional statistical multiplexing gain leads to 15-16 HD channels in 24 MHz of RF bandwidth. This translates to 4 HD channels per 6 MHz of RF bandwidth with no degradation, or double the current number of channels. And this doubling would also apply roughly to DOCSIS data downstreams, where users could double current download speeds as an effective counter to competition, or as a means of attracting small to medium businesses to cable modem service.

Note that as mentioned earlier, additional compression technologies such as MPEG AVC (also termed MPEG4 part 10 or ITU H.264) will further improve the bandwidth utilization of digital video. High definition video using AVC is projected to use between 3 and 7 Mbps, depending on the content type, with sports programming being one of the more difficult types. By comparison, current MPEG-2 HD transmissions typically use approximately 18 Mbps for the video stream. Higher data rates can be used with AVC for even better quality, and lower date rates can be used where some degradation is acceptable and for content that is relatively easy to compress. As a result, in general approximately 2.5 to 3 times as many AVC HD video streams can fit into the data rate previously occupied by MPEG-2 HD. Combining this with the previous example of 1024 QAM/quad channels, this would translate into 10-12 HD channels per 6 MHz of RF bandwidth. The notion of HD video on demand becomes quite viable under such scenarios.

METHODS FOR INCREASING UPSTREAM CAPACITY

Higher Order QAM

Higher order QAM can also be used on cable upstreams, which means greater than 64 QAM TDMA or 128 QAM/TCM SCDMA can be transmitted. If for example, 256 QAM TDMA is used on the upstream, up to 33% additional capacity is provided by using 8 bits per symbol instead of 6. And this additional capacity can be provided in a completely compatible manner with existing, legacy cable modems, since the burst nature of upstream transmissions means that higher order QAM transmissions can be mixed with lower order QAM in the same manner as DOCSIS 2.0 transmissions are mixed with DOCSIS 1.x transmissions.

But the upstream must be robust to ingress, impulse, and thermal noise conditions. As it turns out, most CMTS vendors have included some form of proprietary ingress cancellation processing in their designs, and the addition of SCDMA and TCM to the DOCSIS 2.0 specification significantly improves the robustness to impulse noise as well as providing several dB more robustness to thermal noise.

Consequently, higher order QAM turns out to be quite viable even on today’s upstreams. In the next section, field tests of higher order QAM on the upstream, including 256 QAM TDMA, are presented, where the higher order QAM was operated reliably even in the presence of three ingress signals.

More Efficient Small Packets

The burst nature of upstream transmissions leads to variation in the benefit of higher order QAM on the upstream, however. Since longer preambles are typically required when transmitting higher order QAM on the upstream, and the preambles of small packets can be a significant portion of overall packet duration on the upstream, the benefits of higher order QAM on small packets can be less than 33%. Note that on medium and large packets, which account for the majority of bandwidth consumed on upstreams without VOIP service, the preamble is such a small fraction of the packet duration that 256 QAM provide a 33% improvement in bandwidth utilization. But especially for small packets using the conventional TDMA approach, the improvement can be less than 10%.

Since transmitting VOIP packets using compressed voice will significantly increase the number of small packets on the upstream, methods of improving the efficiency of small data and voice packets will be required. Several methods are available as extensions to DOCSIS 2.0. First, using synchronous SCDMA, instead of the TDMA currently in use, which is quasi-synchronous, permits reduction of the preamble of small packets without degradation in robustness. The synchronous mode is required to maintain code orthogonality in SCDMA mode, but has the added benefit of reducing the preamble overhead on small packets. For example, a 20-30% reduction in packet size can be obtained for highly compressed voice packets using synchronous transport, depending on the specific burst profile that is in use.

But the packet payload itself can also be reduced. A technique known as dynamic payload header suppression (DPHS), which extends the current fixed payload header suppression scheme of DOCSIS in a simple manner, can be used to reduce the header of small packets to the point where the packet duration is about a third of the original duration for small data packets such as TCP ACKs. This translates into three times the bandwidth utilization of small packets, and when synchronous operation is added, small packets can be up to 4 times more efficient. Unlike schemes that only address TCP ACK packets, DPHS also applies to other data packets and to voice packets. On a data-only network, where only 12% of the bandwidth is consumed by small TCP ACK packets, DPHS can provide about 12% overall network bandwidth utilization improvement, while a TCP ACK-only technique would only provide about 8% capacity improvement. On a network with say 50% compressed voice and 50% data traffic, the results are 16% improvement for DPHS and 4% for ACK-only techniques, while on an all-voice network, DPHS can provide up to 25% improvement while ACK-only techniques provide no improvement.

Thus by combining higher order QAM with techniques to address small voice and data packets such as synchronous CDMA and DPHS, the overall bandwidth utilization on the upstream can be increased by up to 33% regardless of packet size. Synchronous mode of operation and DPHS have no impact on ingress robustness, and both can be applied to SCDMA mode in order to be robust to impulse noise. The fact that higher order QAM on the upstream is robust to ingress is born out by field tests described in the next section.

HIGH-ORDER QAM FIELD TESTS

1024 QAM Downstream Test

A test of 1024 QAM was performed on a live cable plant in Rogers Cablesystems that was well-maintained, but nonetheless had measurable levels of CTB and CSO. The test setup is shown in Figure 2 at the end of this paper. The 1024 QAM transmitter was located in the headend while the cable modem 1024 QAM receiver was located in a van and was connected to the cable plant in a residential location. There were 3 active components (amplifiers) between the fiber node and the CM 1024 QAM receiver. No degradation to existing adjacent 256 QAM carriers resulted from the test.

First, a check of 256 QAM operation was made using transmit power levels that were identical to those used by current digital transmitters, and no errors in transmitted packets were observed. Next, since the SNR appeared to be sufficient for 1024 QAM, the modulation was increased to 1024 QAM using the same power level. Errors were detected, and thus the transmit power level was increased by 6 dB, and the system rebalanced, with a resulting SNR of about 36 dB. The majority of errors disappeared, however occasional errors were seen at random times which thus could not be ascribed to CTB and/or CSO cresting as described in [4] since they were not periodic. Possible causes include hardware and software issues in the prototype system used.

Further tests of higher order downstream modulation are in process and may be reported at the NCTA National Show.

256 QAM Upstream Test

The same prototype system was used to test high order QAM on the upstream. In this case, a range of upstream frequencies was made available for testing, some of which included up to 3 ingressors. The range was small however, hence lower symbol rates had to be used in the test in order to compare ingress free operation to operation in the presence of ingress. First, 64 QAM operation (the current maximum TDMA mode for DOCSIS 2.0) was validated in RF spectrum that was free of ingress using an 800 kHz wide signal. Next, the signal was moved so that three ingressors were present and using ingress cancellation processing, 64 QAM operated reliably with less than 0.01% packet error rate (PER).

Next, higher order QAM was tested, both with and without ingress present. 128 and 256 QAM TDMA were seen to operate reliably with less than 0.1% PER when no ingress was present. The 256 QAM signal was then moved to a frequency where 3 ingressors were present and with the ingress cancellor disabled. The PER rose to 96%, however when the ingress cancellor was engaged, the PER dropped to less than 1% PER.

CONCLUSIONS

The deployment of high definition TV will challenge cable operators to find new ways to expand their downstream bandwidth. The techniques presented here have the potential to double the number of HD channels, and when combined with emerging video compression schemes, can provide up to three times the number of HDTV channels in a 6 MHz RF downstream channel.

Wider channels can and have been implemented in current silicon for cable technology. The fact that error-free operation could be achieved at 36 dB SNR on real cable plants confirms the viability of 1024 QAM as a downstream modulation technique as well.

On the upstream, higher order modulation, when combined with techniques such as dynamic payload header suppression, provides robust and reliable data return service to residential customers that provides up to 33% improvement in bandwidth utilization, even on plants with ingress present on the upstream. In particular, 256 QAM on the upstream was shown to be ready for deployment in today's cable plants.

REFERENCES

[1] See www.replaytv.com

[2] “System-level capacity and QoS in DOCSIS1.1 upstream,” C. C. Lee and Jérôme Bertorelle, SCTE Emerging Technologies Conference, January, 2002

[3] , “FastChannel: A higher-speed cable data service,” B. N. Desai, Z. Jiang, N. K. Shankaranarayanan, D. Shur, A. Smiljanic, T. J. Totland, J. van der Merwe and S. L. Woodward, SCTE Emerging Technologies Conference, January, 2002

[4] “Statistical properties of composite distortions in HFC systems and their effects on digital channels,” R. D. Katznelson, NCTA Cable 2002, May 2002

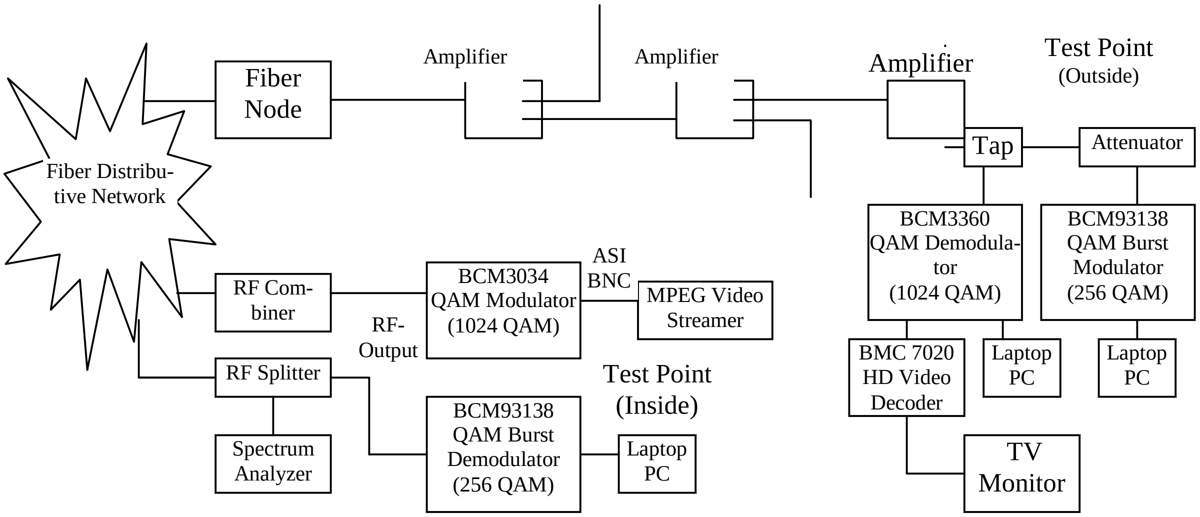

Figure 2. Field Test Setup