Analysis And Prediction Of Peak Data Rates Through DOCSIS Cores

A Technical Paper prepared for SCTE•ISBE by

John Holobinko
Director, Cable Access Business Strategy
Cisco Systems, Inc.
Fort Mill, SC
jholobin@comcast.net

Glenn McGilvray
Manager, Product Marketing
Cisco Systems, Inc.
San Jose, CA
mcgilvg@cisco.com

John Ritchie
Principal Engineer, Engineering
Cisco Systems, Inc.
Lawrenceville, GA
ritchij@cisco.com

.

Table of Contents

Title | **Page Number**

List of Figures

Figure | **Page Number**

Introduction

For nearly as long as DOCSIS has been deployed in cable networks, capacity planning has been a challenge for operators. As the nature of traffic has evolved across the IP network, the accuracy and usefulness of traditional DOCSIS capacity modelling techniques has diminished.

This paper describes an alternative, simpler model for predicting the peak capacity of a DOCSIS network, as measured bidirectionally through the CCAP (converged cable access platform) core, solely based on the number of subscribers attached to the DOCSIS platform and the demographics of the customers. This model is based on analysis of current traffic patterns and traffic types and subscriber behaviors within DOCSIS networks, which is presented herein. Core peak data was sampled over multiple months and across dozens of systems, and is used to show that the accuracy of the technique is superior to other modelling methods.

1. Traditional DOCSIS Peak Capacity Models

Historically, primarily two models have been used to calculate required DOCSIS core network capacity, both dependent on an assumption for the value of statistical multiplexing. Each method was developed based on the results from empirical measurements. Some operators have used a combination of the two models. Both are in limited in their efficacy by the rapidly changing nature of DOCSIS traffic, and therefore misunderstood.

The first model employs a multiplier times the maximum data service speed offered to compute the total capacity required for a service group, then multiplies this times the number of service groups, and finally divides this total by an estimated statistical multiplexing conversion factor.

The second model is bandwidth oversubscription. In the oversubscription model, the maximum data speed of each subscriber within the service group is summed together, then divided by a number Y, called the oversubscription ratio. Oversubscription is an expression of statistical multiplexing. The oversubscription ratio Y:1, depends on the number of subscribers per service group, and more importantly on the types of services that are being offered.

Neither of these models is particularly accurate for predicting the peak traffic rate of a specific CCAP platform.

2. The Evolving Nature of DOCSIS Traffic

From the time DOCSIS was implemented and for the following fifteen years, web browsing and file transfers made up the majority of internet traffic. These data flows are characterized by their short, bursty natures. They do not consistently use bandwidth over a long period of time, i.e. minutes to hours. Based on this burstiness, when looking at aggregate data rates across a network, the aggregate data rate was but a fraction of the total of the data rates offered to all subscribers. This enabled a high amount of statistical multiplexing, i.e. “overselling” the same bandwidth multiple times over. For example if there were 500 subscribers, each with a 10 Mbps (megabits per second) service speed, the computed maximum bandwidth would be 5 Gbps (gigabits per second. However, based on the bursty nature of the traffic, operators found that oversubscription ratios of 50:1 (i.e. 100 Mbps vs. 5 Gbps) were not overly aggressive and provided acceptable performance to this number of data subscribers.

With the advent of Netflix, YouTube and other OTT (over the top) video content providers, over the last 5-7 years the nature of IP traffic has changed immensely. According to the September 2017 Cisco Visual Networking Index[1], globally, IP video traffic will be 82 percent of all consumer Internet traffic by 2021, up from 73 percent in 2016. Global IP video traffic will grow threefold from 2016 to 2021, a CAGR (composite annual growth rate) of 26 percent. Internet video traffic will grow fourfold from 2016 to 2021, a CAGR of 31 percent. Video traffic is very different than file transfers and web browsing in that it creates a highly persistent traffic flow, versus an intermittent bursty traffic flow. Even with adaptive bit rates and modern compression that utilizes chunking and large data buffers, video traffic is highly persistent compared to these other traffic types. This means that the "peak" data rate is sustained over long periods of time, i.e. minutes to hours versus seconds. With the ratio of video to overall internet traffic continuing to expand each year, data persistency is only going to increase. This leaves far less benefit to utilizing statistical multiplexing as a means of bandwidth oversubscription.

Statistical oversubscription calculations depend on the bursty nature of internet traffic. As this burstiness continues to be mitigated by persistent data streams, the accuracy of these calculations decline. Whereas 50:1 and even higher oversubscription ratios were used years ago, every operator continues to gradually reduce their oversubscription ratios each year to a fraction of their former values.

3. Breakdown of the traditional bandwidth models

To use maximum data service speed as a means to compute required capacity, a normal rule of thumb is to take a service group of "N" subscribers. The traditional formula is to take X times the maximum downstream data speed offered where X is usually between 1.5 and 2.0, and use this to compute the number of downstream DOCSIS channels required to support this computed bandwidth. The goal is to have sufficient bandwidth to support at least one user in a service group performing a downstream burst speed test while simultaneously supporting other users' normal data consumption. The required core processing capacity is computed by multiplying this number (the computed bandwidth for each service group) times the number of service groups supported by the CCAP core. The challenge with this formula is that the statistical peak traffic through any service group is dependent upon the number of subscribers in the service group. So, in a large service group with many subscribers, the peak to average will likely be smaller than in a small service group. Secondly, as service groups are summed across a given CCAP core, all peaks will not occur simultaneously. Therefore, computing core capacity using this technique will result in a highly over-engineered core network, while not guarantying adequate bandwidth for each service group.

As maximum data speeds have increased and the nature of IP traffic has changed, the accuracy of previously referenced techniques to compute required network capacity has declined. For example, as data speeds have grown to 1.0Gb/s and even greater, the multiplier "Y" times maximum data speed might remain sufficient for determining the minimum bandwidth required for service group, but it has become a poor method for computing overall required CCAP platform bandwidth, significantly overestimating required bandwidth through the core. Therefore it is instructive to look at the various points of congestion that can occur in a CCAP system.

4. CCAP Congestion Points

In any network, congestion occurs when the total traffic demand exceeds the peak capacity. Any such location is referred to as a bottleneck. In CCAP systems, congestion can occur at the service group level, the line card level, in the switching fabric, in the core, or at the network ingress/egress point. These points are shown in Figure 1.

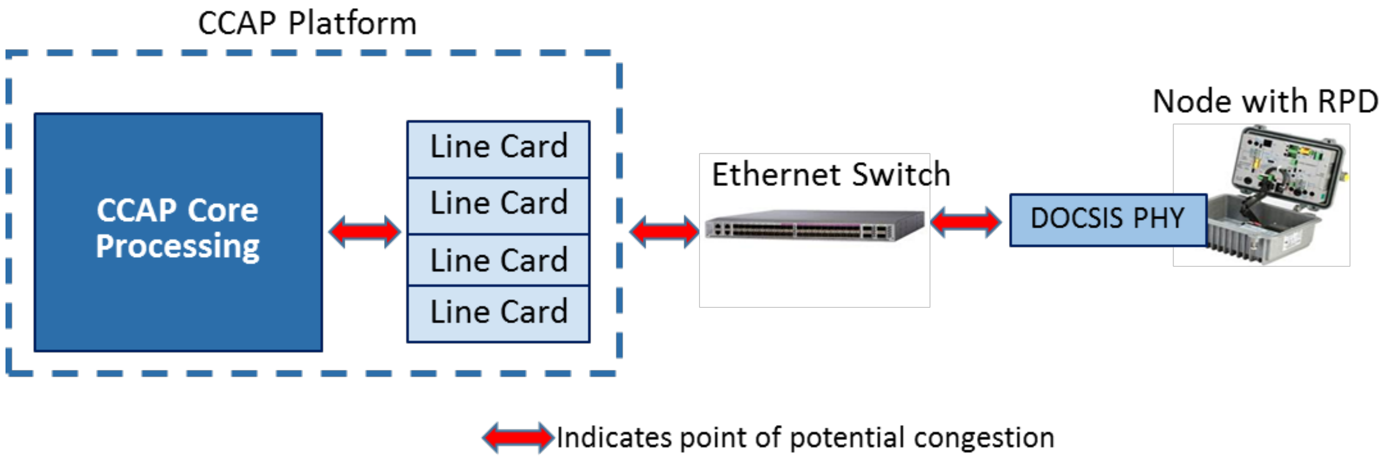

Figure 1 – Potential Congestion Points

5. Service Group Level Congestion

When congestion occurs at the service group level, it can either impact the throughput of all users, or the maximum data speed of the highest speed service tier users. These are two distinctively different forms of congestion. If throughput is impacted across all users, it means that during peak usage periods, the total demand for bandwidth exceeds the bandwidth capacity of the service group. This type of congestion is based on the total number of active users and the type of content they are using. For example, watching high quality videos creates long term persistency of streams, with far less benefit from statistical multiplexing. This congestive impact tends to be over longer periods of time (minutes or hours versus seconds). In contrast, one or more very high speed users may perform speed tests at the exact same time. If insufficient capacity exists, they will not be able to attain their maximum service speed. However, overall network response may not be noticeable to other users depending on the overall QoS settings.

Therefore, when provisioning capacity to a service group, the design goal should be to provide sufficient capacity (i.e. DOCSIS spectrum/bandwidth) to minimize congestion during peak usage periods, while insuring that there is also sufficient bandwidth to support X times (typically X=1.5) the maximum speed of the highest tier of service. Depending on the number of subscribers per service group and their relative internet usage, either the first calculation (large number of subscribers) or the second calculation (fewer subscribers, very high top data speed) will dictate the required downstream capacity.

5.1. CCAP Platform Line Card Congestion

Virtually all dedicated CCAP platforms use one or a combination of an ASIC (application specific integrated circuit) or FPGA (field programmable gate array) to perform framing, scheduling and other functions which are shared across a number of service groups. Typically these functions are done at a line card level or equivalent. At this level, the CCAP architecture should provide sufficient capacity (i.e. DOCSIS spectrum/bandwidth) to minimize congestion during peak usage periods. Since line card speeds are typically 50-100 Gbps, one is not concerned about the peak data speed. Rather the limitation becomes the total number of subscribers, i.e. number of subscribers per service group times the total number of service groups on the line card. In contrast to the service group, the design goal is not to provide sufficient bandwidth, it is to ensure that the maximum bandwidth of the card is sufficient to support the combined peak bandwidth demand for the subscribers connected to the card.

5.2. CCAP Core Congestion

The CCAP core is the point where all traffic passes through the platform in both directions and certain functions such as Quality of Service (QoS) are performed. The goal of the core is to be non-blocking, i.e. support the maximum peak traffic defined by the service levels and QoS. Since CCAP platforms are all architected to last for many years, today the only way to create congestion at the CCAP core is to provision an exceptionally large number of subscribers, or to provide in sufficient network side optical bandwidth to the network core.

6. An Alternative Capacity Planning Model

A 2014 study by Princeton University[1] showed that if cable subscriber currently subscribed service speeds were adequate for daily use (i.e. >25 Mb/s), and these speeds were increased far beyond the service level they were paying for (e.g. to 250 Mb/s) without the subscribers' knowledge, overall data consumption increased by less than 5%. Some subscriber internet behaviors changed, but overall consumption did not. If one started with an assumption that consumption is directly related to data speed, the expected traffic growth should have been 900% instead of 5%. Similar patterns of behavior were noted in Asia when 1Gb/s data speeds were first introduced, i.e. in Asia, operators noticed an initial spike in data traffic by new users for the first month, then a gradual decline until consumption was almost identical to the prior rate of consumption prior to the introduction 1Gb/s data speed subscription service. Other cable operators in North America have corroborated this phenomenon.

We therefore started with the following premises:

- As an alternative bandwidth model, the peak data rate through a CCAP core can be predicted simply by knowing the total number of subscribers attached to the core and the annual data growth rate for the area in which the platform is located
- As long as congestion at the network edge is not affected, total traffic through the core will not be significantly affected by service group size

These are key premises, because they turn upside down the normal way of thinking about capacity planning: Service groups, data rates, oversubscription, etc.

6.1. Core Peak Bandwidth Data Analysis

In addition to Cisco Systems' annual report on growth of traffic across the internet, multiple large cable MSO's publish or otherwise report their year over year growth projections for downstream and upstream DOCSIS traffic across their footprints.

We set about to measure the peak bandwidth consumption across multiple CCAP cores installed in select cable operators with varying numbers of subscribers and service groups on five minute intervals to determine network behaviors and to see if there is another way to predict aggregate bandwidth through the CCAP core. Our goal was answer the following questions:

- Is core CCAP capacity consumption independent of service groups?
- Given a geographic region, can core CCAP capacity be predicted simply by knowing how many subscribers are connected to the core?
- What is the differential between peak traffic times and off peak times? Can this be used to advantage in other areas of cable network operation?

6.2. Sample Universe

Our total sample size covered tens of CCAP platforms serving multiple hundreds of thousands of DOCSIS subscribers. Taking samples of the core data rate on 5 minute intervals, the total number of samples was in the hundreds of thousands. The total sample period was seven months. Therefore, we feel very confident that the statistical sample size is well beyond the minimum sample size required to achieve confidence in the results.

So as not to reveal any cable operator sensitive data, we will generally summarize the data sets as follows:

- The total sample size consisted of approximately 40 CCAP platforms. Some had converged MPEG video and data/voice services, while others were data/voice only.
- Demographics ranged from dense urban to suburban to rural areas.
- The number of subscribers per platform varied by over 500% with the largest number being approximately 12,000 subscribers. The total number of subscribers across all sampled systems was approximately 300,000.
- The number of service groups per CCAP platform varied by nearly 500%, i.e. the fullest platforms supported four times the number of service groups as the number of service groups the least utilized platforms supported.
- The number of subscribers per service group also varied by over 400% as well with the largest service groups having more than 300 subscribers attached.
- On each CCAP platform core, a sample of the total combined forward and return path combined throughput was taken on five minute intervals. The total sample period for the study was seven months. The total samples for the study approached one half million.

Figure 2 shows the peak data rate through CCAP platform cores over the sample period. If the peak bandwidth were proportional to the number of DOCSIS channels provisioned per service group, the most variation we would expect is 4:1, based on the smallest platforms having one quarter of the service groups as the largest platforms. But instead, we see a ratio of more than 6:1.

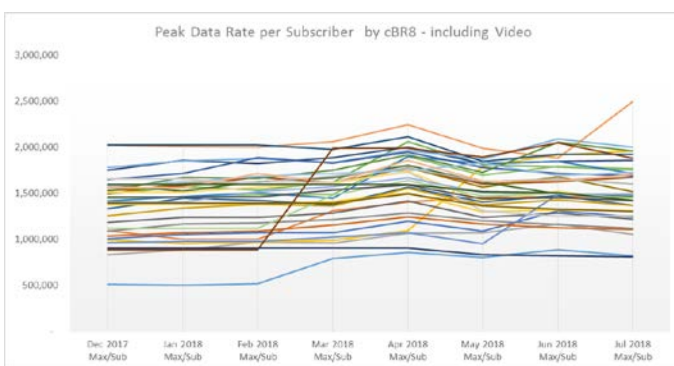

Figure 2 – Peak Data Rate per CCAP Platform

This can neither be explained away by converged video on only some of the platforms. Figure 3 shows the same data with video subtracted from calculations; only the data traffic on each platform is presented, whether or not MPEG video/data convergence was present on a platform. The ratio of the highest peak traffic platforms to the lowest peak data platforms remains the same.

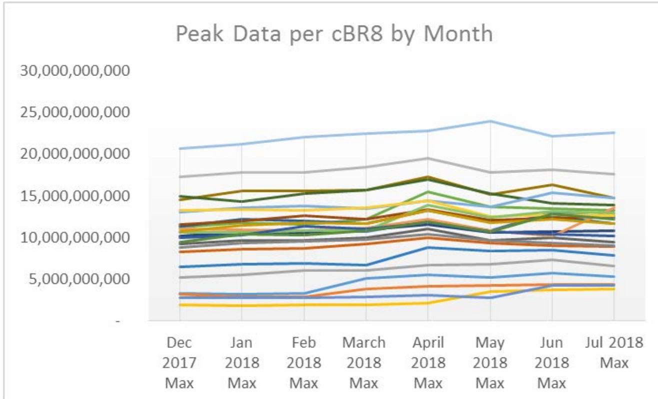

Figure 3 - Peak Data Only Rate per CCAP Platform by Month

However, the data becomes much closer if it is normalized to a peak average per subscriber rather than the individual CCAP system peak throughput itself. This number represents the average bandwidth per subscriber at the point where the platform reaches its absolute maximum throughput (i.e. customer demand) in any given month. This peak average per subscriber is calculated for each platform by taking the monthly peak traffic sample and dividing it by the number of subscribers.

Avg Peak Traffic Rate per Subscriber (APS) = Platform Peak Traffic Rate/Subscribers on Platform

Figure 4 shows the resulting graph comparing the APS across all of the platforms, (no video). The variance across all of the machines is now far closer, 2:1. The highest APS is approximately 2.2 Mbps while the lowest is approximately 1.1 Mbps. (There were a small number of systems that were activated during the study. These are shown with the very high growth rates). Looking more closely at the data reveals that the systems with lower APS tended to be in more rural areas, while more dense metropolitan areas generated higher traffic demand. What the data shows is that normalizing the peak data on a per subscriber basis is a far more accurate predictor of total peak bandwidth through a CCAP platform than using other modelling techniques such as bandwidth per service group along with a multiplexing factor, or total provisioned bandwidth across the platform, number of service groups, etc.

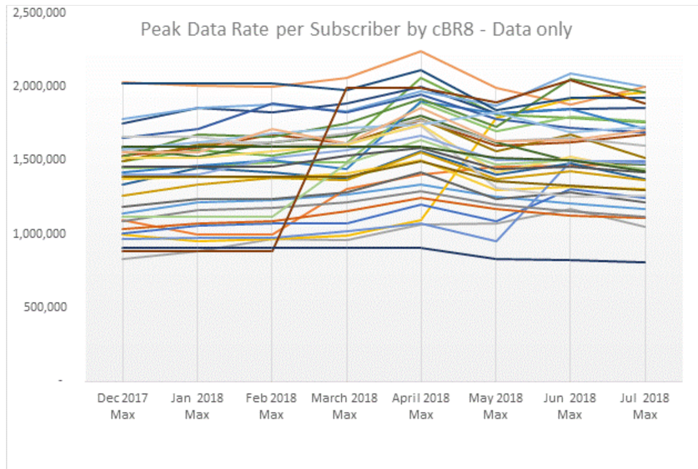

.

Figure 4 - Peak Data Rate per Subscriber by CCAP Platform – Data Only

The value of this is in the ability to predict CCAP platform total peak bandwidth in the future. The data shows that by knowing the number of subscribers attached to a given platform and whether it is rural or urban, one can make a good estimate of future bandwidth requirements. For example, in an urban environment, we would predict that subscriber APS is conservatively 2.5 Mbps or less as of the date of this paper (recognizing that data grows annually, thus affecting the baseline). Figure 5 below shows graphs based on 30%, 40% and 50% composite data growth for systems serving 5,000, 10,000, and 15,000 subscribers respectively. 2.5 Mbps per subscriber is the starting baseline for these graphs. One could do the same exercise for rural systems, or simply superimpose half the urban APS for rural systems.

It should be noted that these results are for North America. In other regions where viewing habits and internet usage varies dramatically from North America, a different baseline and growth rate may be necessary. However, the fundamental model will be the same.

Figure 5 – Predicted Peak Data Demand Based on Subscribers and Growth Rate

Note that as long as there is not congestion within the bandwidth of the SG's, the number of SGs that are served by the CCAP core do not impact throughput, only the total number of subscribers. This means that modelling the total peak traffic across a disaggregated core with many small SGs becomes very simple. As long as there is sufficient bandwidth in the SG to support the peak bandwidth demand of the SG (i.e. no bandwidth bottleneck), then the results will be the same for a given number of subscribers whether they are in X service groups or divided amongst 4X service groups.

7. More Observations On The Nature of CCAP Core Traffic

If we were to sample the same system for more than a year we would expect the peaks in the second year to be N% greater than the peaks in the prior year. This percentage N would be equal to the annual data growth across the network. However, on a month to month basis traffic is not necessarily N/12, but can vary significantly either up or down from this %, and in some cases be actually negative for a given month. Therefore, all of the graphs in this study have an underlying annual growth rate within the data. We could not adjust monthly traffic apart from the growth rate by subtracting a growth line without having two years of data from which to extract the growth rate. Perhaps in another year, a follow on paper can address this.

7.1. Peak to Valley Traffic Analysis

Figure 6 shows the peak to valley data rate for each platform, on a monthly calculation basis. For some systems, the peak data rate is more than 12 times the valley, while for the least variable systems the peak to valley is approximately 4:1.

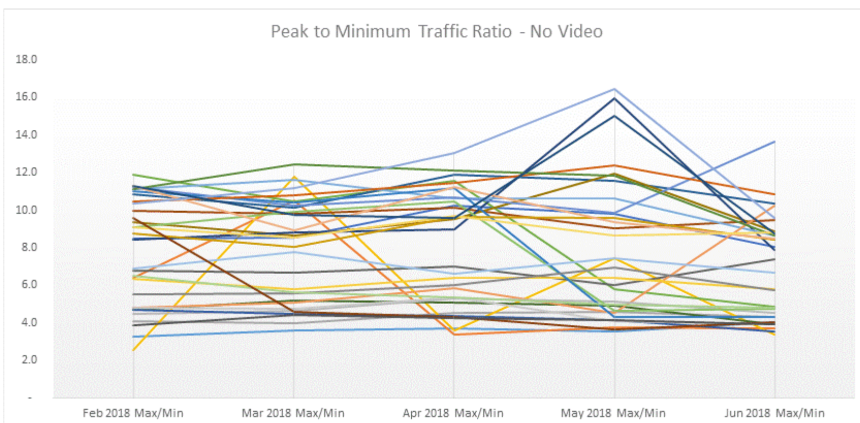

Figure 6 – Peak to Minimum Traffic Ratio Across CCAP Platforms

Figure 7 shows the monthly samples for one system in the study. Note on a daily basis the daily peaks to valleys.

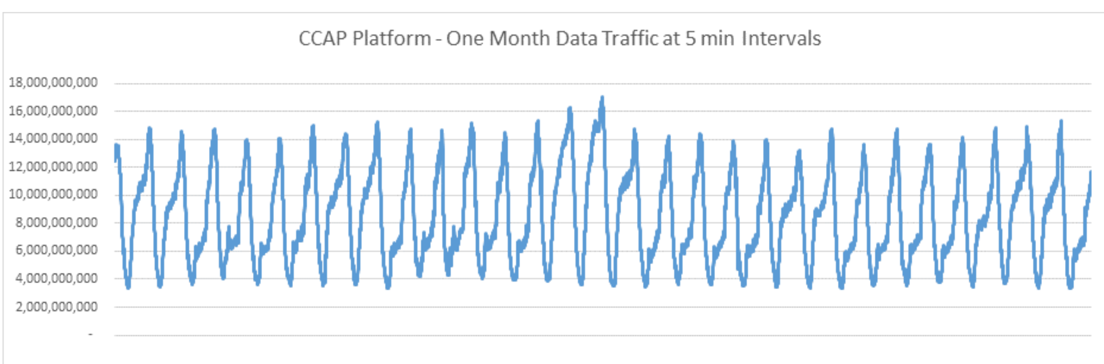

Figure 7 - Single CCAP Platform Traffic for One Month Taken On 5 Minute Intervals

Examining one day in detail, one can see the significant difference in traffic as shown in Figure 8.

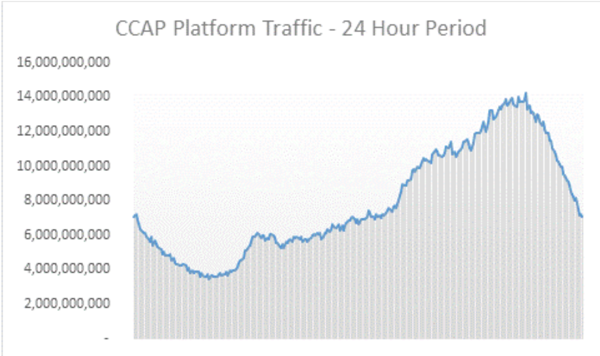

Figure 8 – Single CCAP Platform Traffic for One Day Taken On 5 Minute Intervals

What can be garnered from this information is that the differences in traffic are significant based on time to day, and further, appear to be well behaved. This may have significant consequences in the future relative to the ability to manage cloud based processing resources and to manage power consumption across the network, gaining appreciable savings by lowering the power required to support low traffic time periods.

There have been previous suggestions to turn off DOCSIS QAM carriers during lower usage times as a means to save power in RF amplifiers. This is illustrated in Figure 9. However, turning off DOCSIS carriers turns out to be impractical. In D3.0 and D3.1 the channels in a bonding group are fixed and are programmed into each cable modem in the SG. If a channel is no longer present, the modems interpret this as a failure and start communicating this back to the CCAP core as a problem. The alternative is to redefine the SG at different points of the day. But to do this requires re-provisioning every modem, taking them offline during the process, resulting in daily service outages.

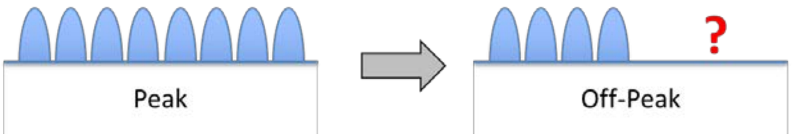

Figure 9 – Proposed Power Savings By Turning Off Carriers During Low Demand Periods

An alternative potential means to achieve power savings is to reduce the modulation profile and lower the corresponding carrier RF level of each DOCSIS channel accordingly. With regards to modulation throughput, 4096 (2^12) QAM (Quadrature Amplitude Modulation) provides double the transmission speed per bit as 64 QAM (2^6). However it takes 18dB more power (translating to approximately 60 times the power) to transmit a 4096 QAM carrier than a 64 QAM carrier. This additional power is a significant operations expense.

Referring again to Figure 7, the peak to valley speed in a single day on this illustration varies by 4X, and as seen from Figure 5, can vary by up to 12X on some systems. However, if we were able to set the highest power level safely above the peak data speed for the month and only reduce the power such that the minimum speed we supported was half of that, this translates to an 18dB power savings at off peak times. We can set the peak power to support a data rate higher than the peak, to enable a safety zone, and a low level significantly higher than the lowest data rate or highest data speed offered.

For example, if our DOCSIS spectrum supports a maximum data throughput of 5.0 Gbps at 4096 QAM, then at 64 QAM the same spectrum supports 2.5 Gbps. We could safely save 10 dB of power across the DOCSIS spectrum, and achieve a minimum data rate of approximately 3.0 Gbps.

Superimposing this on the CCAP platform one day data shows how this might be employed. This is shown in Figure 10. In this example, we use only 10dB for the difference of the signal levels for reasons described below. If the peak rate that can be supported during low traffic time periods is 3.0 Gbps, then we can comfortably have a minimum supported service speed offering of 2 Gbps and still achieve 90% output power savings for that part of the spectrum during low usage time periods.

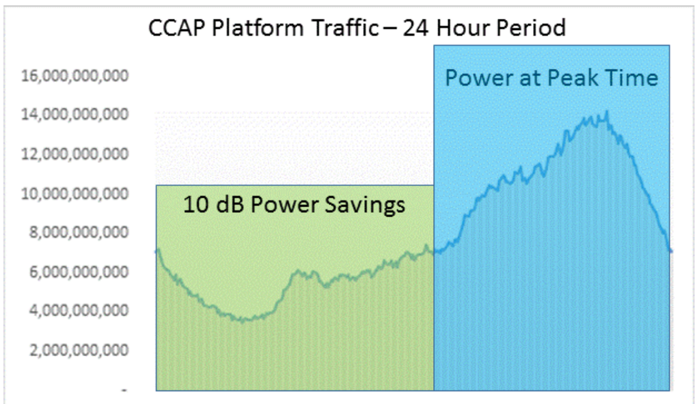

Figure 10 – Lowering QAM Modulation Profiles During Off Peak Periods

In the CableLabs RPD (remote phy device) spec, there is an adjustable output of 10dB for every carrier. And each one can have its modulation profile changed. So, for all DOCSIS carriers, this means that their levels can be reduced in off-peak periods by up to -10dB which is still highly appreciable (-18dB not possible based on adjustment range). This is depicted in Figure 10.

In an all-IP DOCSIS network, this would result in reducing the required power output of the system hybrids by 10 dB equivalent to approximately 90%. However in a system with MPEG video occupying 50% of the downstream bandwidth, the savings would be only 7-8 dB across the spectrum based on where the DOCSIS carriers are located. In either scenario, since systems never operate at full capacity, we can potentially expect even higher power savings. In today's super high output 1.2 GHz nodes, each RF hybrid consumes over 15 watts at full power. By cutting the composite power even by 6-8 dB will likely result in savings of 7-10 watts per hybrid, meaning saving 28 – 40 watts per node during off peak periods.

Therefore, power savings can become an additional driver to migrate cable systems from MPEG video to full spectrum DOCSIS with all- IP video delivery in the forward path.

But how to accomplish this and not end up in a situation where there is insufficient power and modulation profile to support traffic bursts? The answer is to use analytics. Begin by sampling the downstream traffic on each service group every five minutes. Set the power level and modulation profiles on the DOCSIS carriers such that there is 50% headroom above that level. Compute the appropriate bias voltage/current on the RF hybrids in nodes and amplifiers to support the total composite RF level required. Of course, this requires smart nodes with dynamic remote capability to adjust their power levels. A block diagram of these functions is below in Figure 11.

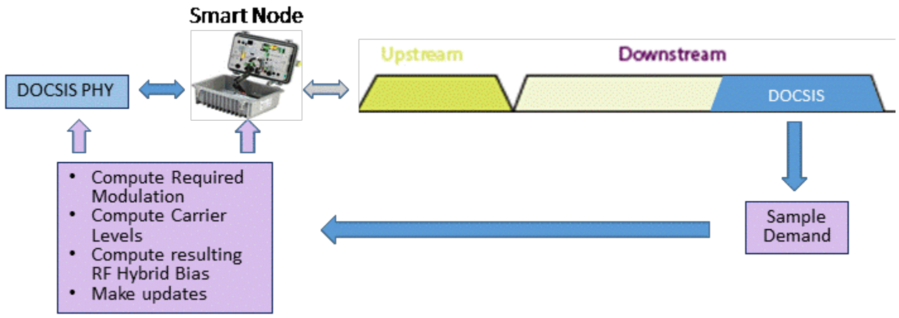

Figure 11 – A Method For Dynamic Power Savings Using Analytics and Smart Nodes

Cisco has applied for a patent that covers this process.

8. Further work

This study is a first step in big data analytics application for traffic prediction. A next step is to take the same sample interval as used in this study and to apply it to measure the peak traffic for every individual service group within a single CCAP platform. One would expect that the larger the average size of the service group the more that individual SG peaks and valleys will reflect the core peak valleys (after all, for a CCAP platform with 50 SGs, one SG represents a sample size of 2%). However as service groups become smaller, either through having more SGs on the same core, or because the CCAP platform is in a rural area serving far fewer total customers, one would expect the variations between an individual service group and the total to be greater. These assumptions have not been validated with data and are an area for future study.

9. Conclusions

We have demonstrated that using the number of subscribers attached to a CCAP core is a better method for predicting traffic through the CCAP core than by any measure employing service group counts, channel counts and an oversubscription factor. We expect the peak data consumption per subscriber to be very consistent across CCAP platforms with similar demographics, such as mostly urban or mostly rural settings. Geographically, we expect a different number per subscriber in each region (e.g. LATAM, Western Europe, and Asia) based on the cultural differences that are reflected in overall internet usage.

One provocative issue that arises from this study is the usefulness of traditional DOCSIS CCAP licensing models which charge for bandwidth on a per service group, or per channel basis. Given that operators revenues are per subscriber and that peak bandwidth is most dependent on the number of subscribers attached to the platform, a per subscriber license plan that enables unlimited bandwidth consumption appears to be the most practical licensing solution. This will become even more important with DAA (distributed access architecture) networks and cloud based CCAP core functions, wherein service groups are no longer an architectural limitation as imposed by dedicated CCAP hardware.

Bibliography

1. Cisco Visual Networking Index: Forecast and Methodology, 2016–2021, Cisco White paper, June 6, 2017

2. Grover S., Ensafi R., Feamster N. (2016) A Case Study of Traffic Demand Response to Broadband Service-Plan Upgrades. In: Karagiannis T., Dimitropoulos X. (eds) Passive and Active Measurement. PAM 2016. Lecture Notes in Computer Science, vol 9631. Springer, Cham