MOTOROLA

OPTIMIZING DOCSIS UPSTREAM UTILIZATION – CAPEX AND OPEX

MARCH 2013

TABLE OF CONTENTS

MOTOROLA

INTRODUCTION

In 2008, Motorola discussed how to get "Better Returns from the Return Path: Implementing an Economical Migration Path for Increasing Upstream Capacity." That paper can be read in full here. This paper expands on the subject by analyzing how new DOCSIS 3.0 upstream CMTS technology offers better opportunities to optimize the ever-evolving HFC network to satisfy the needs of this increasingly vital business — high speed internet (HSI).

Over the last few years, MSOs have begun a public transformation from marketing themselves as video purveyors to "Broadband" operators. New companies providing video-over-the-top, social networking and on-line gaming have grown so rapidly and become so sticky to consumers that the HSI service has become the centerpiece of the triple play offering. While the cable operator has fared better than most in capitalizing on this trend by growing its share of broadband subscribers, keeping pace with traffic growth continues to require large investments in HSI networks. The market may allow for HSI average revenue per user (ARPU) to climb slightly, but the rate of traffic growth demands on the network and the investment required to keep pace threaten to dramatically compress operating margins if MSOs continue to pursue modus operandi of network design. This is particularly true with how MSOs address the so-called Achilles Heel of the HFC network: that precious 5-42 MHz of upstream bandwidth available in most cable networks.

In the early days of DOCSIS 1.0 deployments, it was common to have HFC networks with 2,000 Homes Passed (HP) per Fiber Node (FN). Since penetration (subscribers per FN) was so low, and 56 kbps dial-up modem service was the competitive offering, a common Cable Modem Termination System (CMTS) design would be one downstream port offering 35 Mbps serving eight FNs. This is commonly referred to as the Serving Group (SG). Each of these FNs would have one CMTS upstream port, typically using QPSK modulation on a 1.6MHz channel, which provided about a 2 Mbps pipe. A common CMTS RF module design was therefore a 1x8 (see figure 1).

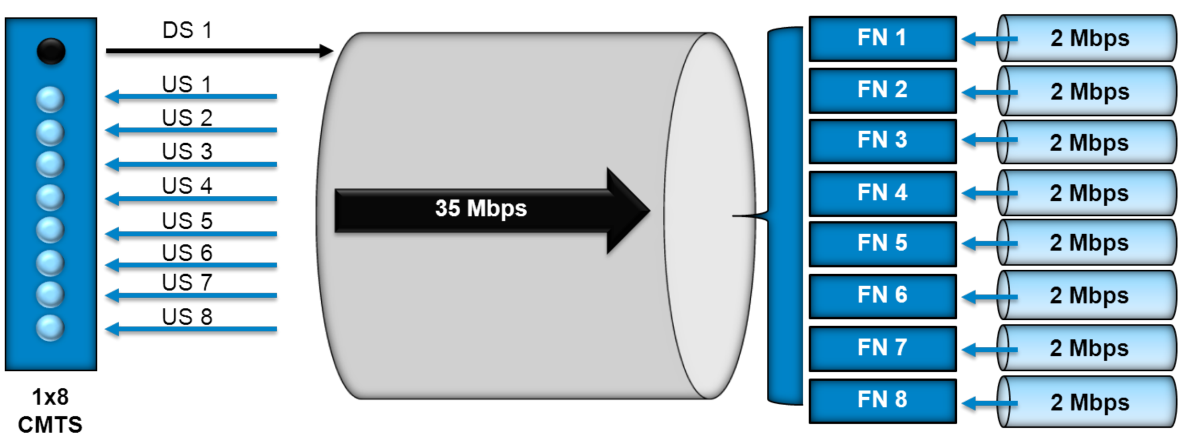

Figure 1. 1x8 CMTS Architecture (one downstream port and eight upstream ports)

As more subscribers signed on to the network and competition and demand for bandwidth grew, the design of the network evolved; the 1x8 became a 1x4 or 2x8 module. Through this evolution, it was common practice to match the slower pipes from upstream ports with one FN and to combine FNs onto the faster downstream ports. This tried and true method of traffic management worked well when coupled with the practice of node splitting. When a downstream pipe filled up, the FN was split in half and a new CMTS port was deployed. This led to a sort of linear expense where the CMTS budget began consuming a growing portion of the MSO's capital expense budget.

The need to offer faster internet speeds and to deal with bandwidth management more cost effectively than what was available in DOCSIS 2.0 led to the creation of a new specification. DOCSIS 3.0 enabled more available bandwidth through bonding of multiple channels to create bigger, faster pipes. Now, the persistent traffic problem of downstream congestion (remember those super jittery video clips?) could be solved by delivering more channels or pipes to the same SG rather than having to split the FNs. Today, it is very common to have four downstream channels from one CMTS port dedicated to several FNs. Many operators have plans to move to eight in the next year or two. The driving force behind DOCSIS network evolution is the trend toward greater amounts of OTT video consumption over the HSI service. In order to cost effectively keep pace, CMTS vendors are leveraging silicon designs that enable increases in the number of downstream QAMs per port. CMTS vendors are showing new technologies with 8-64 QAMs per DS port in their current CMTS and next generation Converged Cable Access Platform (CCAP) products.

So if the answer to cost effectively addressing the downstream traffic demand is greater QAMs and more FNs per port, what is the plan for the upstream? What advances have been made in DOCSIS upstream technology to deal with the much noisier segment of the spectrum to keep pace with the new demands on the network? How do we optimize the implementation of new DOCSIS 3.0 CMTS upstream cards, which after all are as costly as the downstream? This paper intends to address each of these questions in detail to help define the proverbial "sweet-spot" to satisfying the diverging interests of the MSO's customers and its owners or investors.

CATV NETWORK UTILIZATION

It is important to design and maintain CATV networks to support variation in performance that could be influenced by factors including temperature changes, level variation over frequency, and impairments. An optimized CATV network will have adequate performance margin to accommodate these changes, all while delivering services that customers desire. If there is too little margin, then high occurrences of network failures will result, while too much margin indicates an over-engineered design and wasted investment dollars. As an example, many cable operators will maintain a CATV network with performance margin in the form of 3-6 dB more Signal-to-Noise Ratio (SNR) than what would be the minimum required to support a given modulation profile (higher modulation provides greater throughput but requires higher SNR). This performance margin represents the sweet-spot between objectionable occurrence of failure and reduced return on investment (ROI).

A similar sweet-spot exists for network utilization. Network utilization is another aspect of CATV network performance that represents how well CATV networks are used. The utilization we are referring to in this discussion, is the amount of bandwidth capacity that users consume in a given upstream serving group over the amount that is provided. As observed with performance margin, optimized network design and maintenance will exploit the sweet-spot that exists between objectionable failure and reduced ROI. Network utilization is similar to performance margin in that if network utilization is too high, blocking or network failures could occur at an intolerable rate, and if it is too low, it could indicate an over-engineered network. An over-engineered network will result in lower than optimal ROI but can also mask potential performance issues in the network. If network utilization is too low, there will be too few users to provide feedback on network performance, allowing network problems to go unnoticed longer than if there were more users. A common impairment that often goes undiagnosed in a multi-channel upstream network design is intermittent laser clipping. MSOs will design and maintain RF levels at the optimum performance and dynamic range, but in reality this performance is diminished by additional dynamic range consumed from low frequency ingress or non-paying signals also present in the return band. An under-utilized network may rarely be in a state where both the paying and non-paying signals are at maximum power, but when this does occur, brief periods of degraded performance, due to laser clipping, will inevitably result. The hardest part of diagnosing these intermittent problems is catching them when they happen and identifying them. A network optimized for utilization will increase the likelihood of diagnosing intermittent problems because the network will be passing traffic most of the time. Designing and maintaining the CATV network that exploits the sweet-spot of network utilization through a sound headend node combining strategy helps ensure optimized utilization.

UPSTREAM CAPACITY MANAGEMENT

Historically, upstream congestion has been less of an issue than downstream congestion. Put another way, years of splitting nodes and segmenting the network to reduce narrowcast serving groups has been driven by over utilization of the downstream and not the upstream. Fiber nodes are now deeper in the network with far fewer amplifier cascades; naturally, subscribers per node have been steadily decreasing even as penetration grows. This network evolution and other HFC technology improvements has made possible the use of higher order modulation and wider channels in the upstream, yielding greater capacity throughput. This, coupled with higher density, better performing CMTS cards, has opened the door to much more effective ways to manage upstream capacity. One way an MSO can better manage upstream capacity utilization, for example, is with a sound headend node combining strategy. Instead of the lower utilization scenario, where one node upstream is connected to one upstream CMTS port (1:1) as shown in Figure 2, multiple nodes may be connected to one CMTS port. One historic concern of headend node combining is reduced performance and RF level into the CMTS port. Fortunately with today's advanced RF capabilities, such as those on Motorola's BSR 64000 RX48 upstream card, MSOs have a solution to improve dynamic range where fidelity requirements can be maintained with RF input levels as low as –7 dBmV per channel, all while supporting 4 bonded channels on a single port. The most significant advantage of headend node combining is the ability to manage and optimize network utilization and leverage popular DOCSIS features such as modulation order, bonding, multiple access, and FEC to deliver on capacity requirements, while maintaining adequate performance margin.

Because of its advanced RF capabilities, the RX48 on the BSR 64000 enables the MSO to combine nodes and still operate with adequate performance using high order modulation profiles or to maximize RF performance measures and increase upstream channel utilization using lower order modulation profiles. Increasing upstream utilization will enable better "tuning" of the network because network impairments can more easily be identified when there is constant traffic. As an example, consider two candidate 1:1 nodes that are optimized to support 64-QAM with adequate performance margin, but have extremely low utilization. 2:1 combining with no modulation profile changes could reduce performance margin by 3 dB. However, 2:1 combining using 32-QAM instead would offer similar performance margin as the 1:1 scenario. Similarly 4:1 combining using 16-QAM instead would offer similar performance margin as the 1:1 scenario. Each increase in this combining scenario increases the CMTS utilization, due to more subscribers per port and reduced bit rate. The challenge for MSOs is aligning the capacity with utilization over time to ensure their customers get services they deserve while network resources are optimally used.

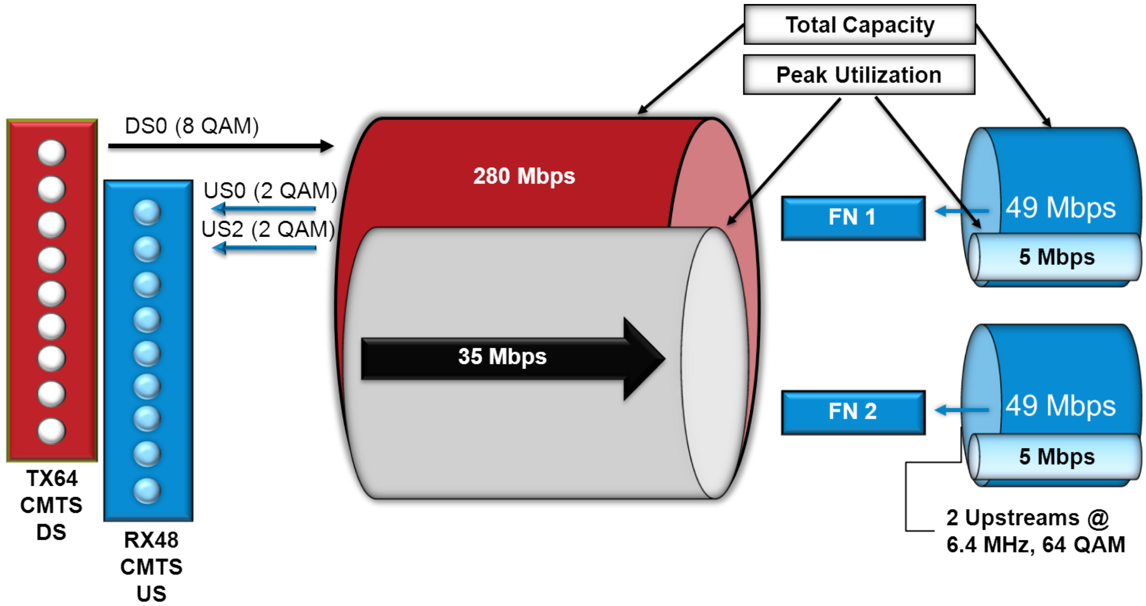

Figure 2. Underutilized 1:1 Headend Architecture Scenario

Two underutilized 1:1 nodes, whose 64-QAM links are used 10% of the time, could combine 2:1 to increase utilization by 20% or 4:1 to utilization by 40%. Additionally, reducing bit rate to restore performance margin would increase utilization estimates as well. Figure 3 illustrates a 2:1 combining scenario where utilization is 24% with equivalent 1:1 performance margin.

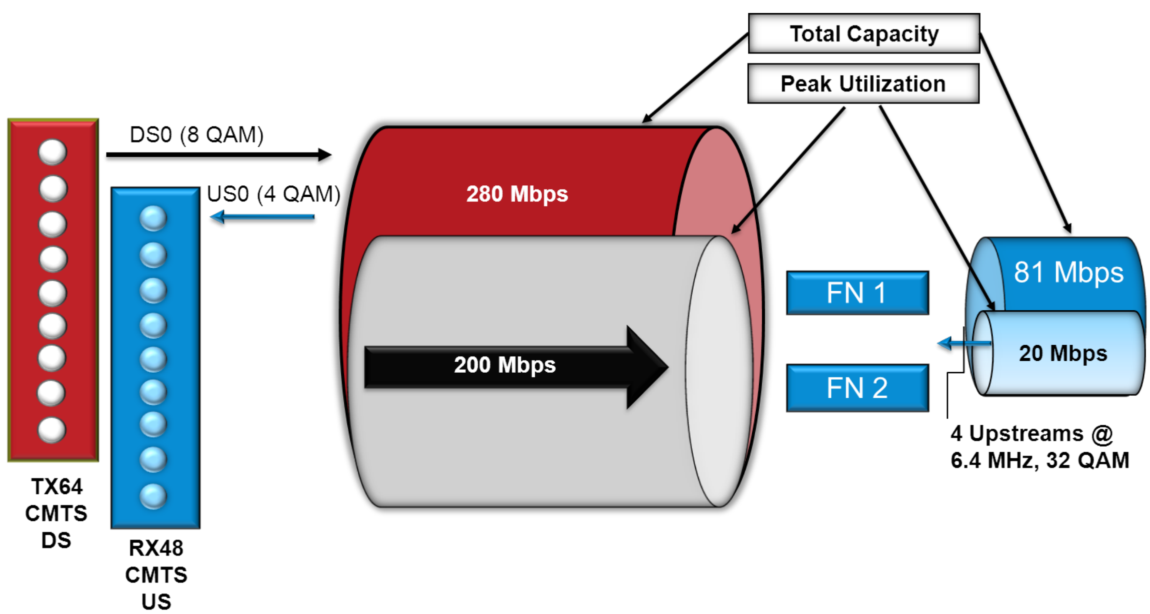

Figure 3. Optimized 2:1 Headend Combining With 1:1 Equivalent Performance Margin

Note that 4:1 combining scenario utilization would be 60% with equivalent performance margin, as shown in Figure 4.

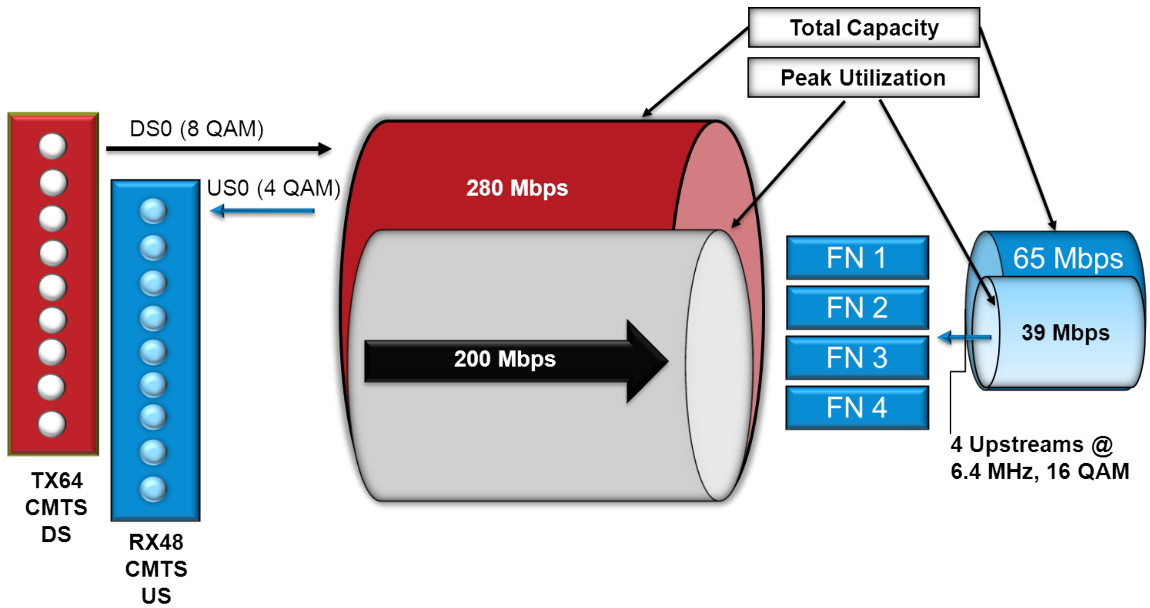

Figure 4. Optimized 4:1 Headend Combining With 1:1 Equivalent Performance Margin

If and when customer use increases over time, capacity can be managed with additional profile and configuration enhancements like bonding additional channels or reductions in headend node combining. The important thing is that these adjustments are executed as demand requires, and resource dollars are only spent when they are needed. A sound headend node combining strategy offers a simple solution, making it easy for MSOs to optimize network design, and make future investments only when upstream network needs justify them. In summary, it is more prudent to spread the capital expense of a CMTS port over a larger serving area, especially when the design meets performance margin goals and increases utilization to levels that provide meaningful network feedback.

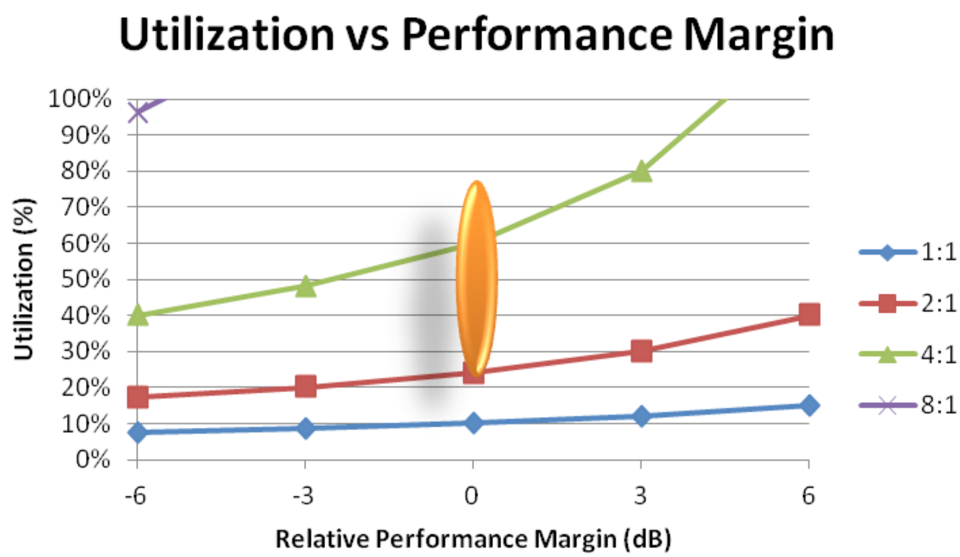

Figure 5. Utilization vs. Performance Margin. This graph looks at the affect of combining nodes on both performance margin and Utilization. The highlighted area represents the "sweet spot" in terms of achieving acceptable performance margin, utilization and Return on Investment.

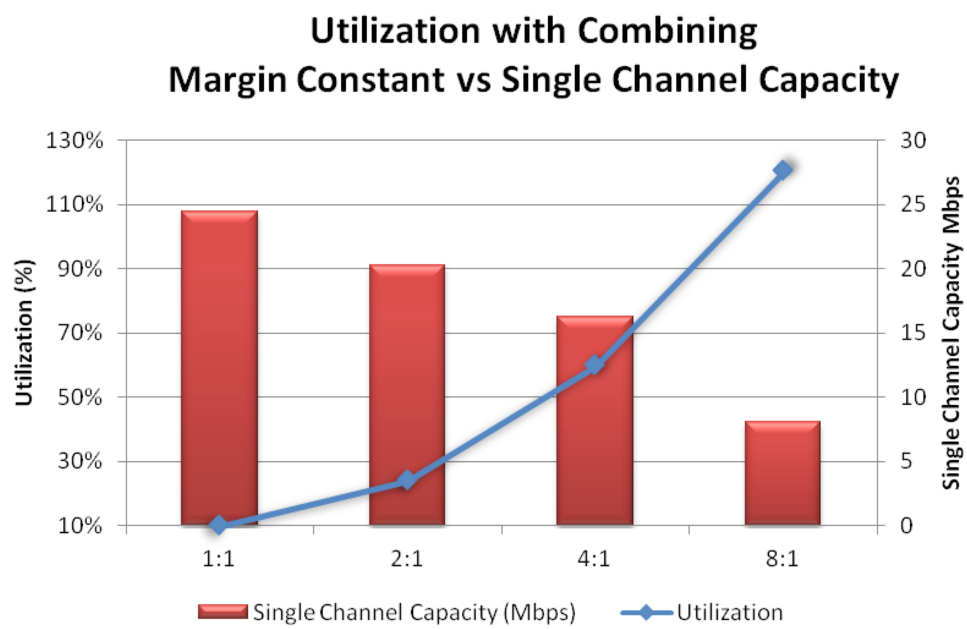

Figure 6. Utilization with Combining Margin Constant at 6 dB vs. Single Channel Capacity. This graph looks at combining nodes and its affect on utilization and single channel capacity while holding performance margin constant. This highlights the trade-off in the speed that the cable operators may offer for any single channel (Non DOCSIS 3.0) customer. Note that capacity available with 2:1 or even 4:1 combining is well in excess of most consumer offerings today. DOCSIS 3.0 upstream channel bonding, would allow increase peak speeds.

These graphs assume the initial system condition is 1:1 and maximum utilization is 10%, with a modulation profile based on 64-QAM. All nodes serve approximately the same number of subscribers, so combining results in doubling utilization. Channel bandwidth is 6.4 MHz wide so that theoretical throughput can be approximated based on ATDMA using concatenation, which allows multiple upstream data packet per upstream burst and request piggy-backing.

Figure 7 expands the above concepts further by accounting for the affect bonded operation has on capacity, while illustrating the role headend node combining could play in creating a suite of sweet-spot options for MSOs. Headend node combining clearly creates options for MSOs even under the most modest set of assumptions. Consider managing capacity demand increase given a starting condition of single channel operation with 8:1 combining and 80% utilization. In this scenario, accommodations for increased capacity can come in the form of increased modulation complexity (assuming the network can support it), reduction to 4:1 combining, or additional bonded channels, as illustrated in Figure 7 by the three respective curves.

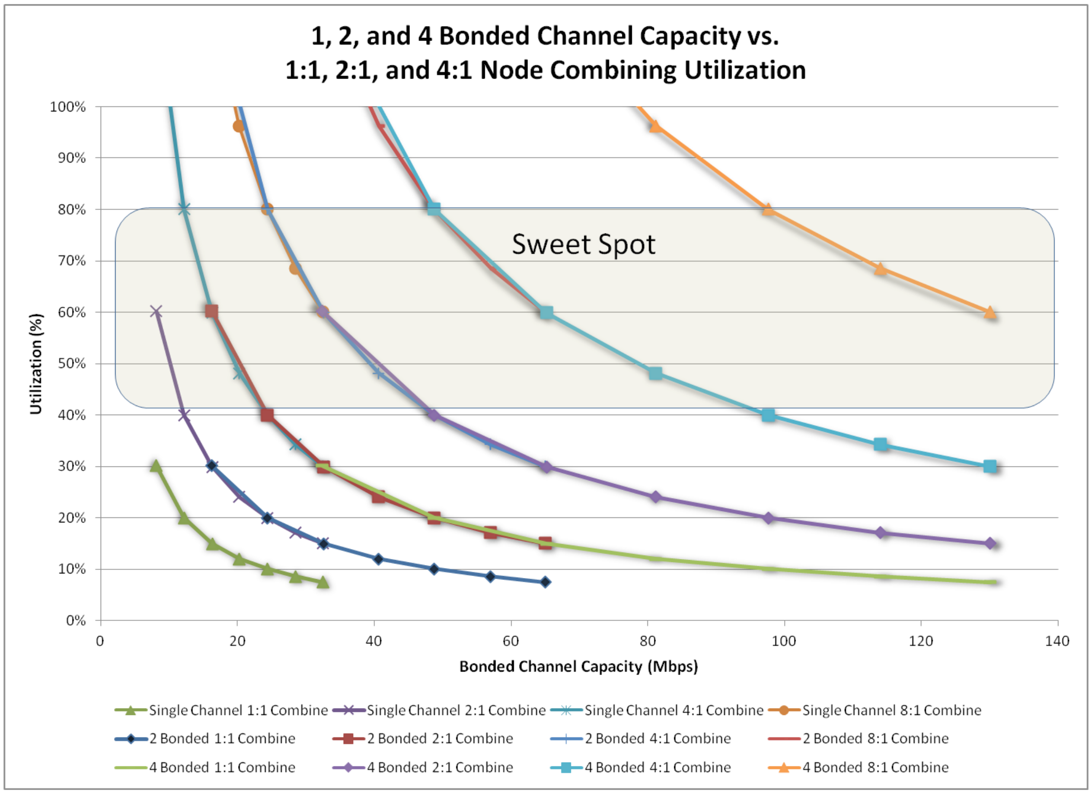

Figure 7.

BENEFITS OF DIAGNOSTIC MAINTENANCE

Back before powerful proactive monitoring tools were developed, earlier headend node combining strategies may have seemed too complex or costly to maintain. Fortunately, due to the CableLabs InGeNeOs® working group coupled with advances in monitoring technologies, networks have evolved tremendously, which helps MSOs to identify, isolate, and mitigate HFC impairment issues quickly.

Further, an RX48 card already supporting 4 bonded channels on a single port still has two additional receivers that can be dedicated to support significant diagnostics, including spectrum analysis, which can be performed easily by the MSO to further enhance monitoring capability. Fast Fourier Transform (FFT) captures can be executed frequently enough to detect spurious energy above the system's diplex filter cutoff frequency, a clear indicator that a subset of the combined optical links are impaired by laser clipping. This diagnosis can be further validated through ongoing monitoring of aligned impact to InGeNeOs® parameters, including pre and post equalization, codeword error rate (CER), and modulation error ratio (MER).

An appropriate mitigation strategy for laser clipping would be to re-align all optical links to account for the additional power from non-paying signals in both the optical transmitter and CMTS input. This fix would have two benefits: first, reducing the relative level of the non-paying signals by the attenuation introduced at the optical transmitter; second, all modem transmit levels will remain the same pre and post execution of the upstream optical link re-alignment.

CONCLUSION

The cable industry has put forth a significant amount of analysis and thought about how to evolve the CMTS to cost effectively address the rising demand for downstream services driven by OTT traffic. The solution is largely understood to be simply greater QAM port densities combining FNs into larger, much faster downstream pipes. Through advances in CMTS RF technology, a similar practice is now able to be applied to economically solving the less talked about upstream capacity challenge. MSOs will benefit greatly from well-executed headend combining strategies that are now possible with modern day DOCSIS 3.0 Upstream technology and enhanced DOCSIS link monitoring capability, with the most important benefit being Capex and Opex savings from right-sizing and optimally using valuable network investments. This white paper has addressed many pros associated with headend node combining as well as how to mitigate impairments such as laser clipping in a node combined environment. Motorola Mobility's BSR 64000 with RX48 DOCSIS 3.0 Upstream Module, with its enhanced dynamic range, is designed to enable MSOs to easily overcome these impairments, dispel previously established consequences of legacy equipment, and find the sweet-spot between performance margin and network utilization, as well as, between customers and investors.

MOTOROLA