WHITE PAPER

Motorola's Leveraging S-CDMA
for Cost Efficient Upstream Capacity

Table of Contents

Motorola's Leveraging S-CDMA
for Cost Efficient Upstream Capacity

Abstract

This paper discusses how S-CDMA provides an efficient means to increase HFC upstream capacity. We start with a brief look at factors driving demand for upstream bandwidth and some bandwidth models to estimate how much upstream bandwidth we are going to need over the next 5-7 years.

From here, we provide a technology overview of S-CDMA and then dive into how to get the most from your S-CDMA system. S-CDMA implementations have not sat still since the DOCSIS 2.0 spec was published. The paper will discuss improvements made with DOCSIS 3.0 as well as continued support of DOCSIS 1.1 modems using the concept of logical channels. Critical factors in upstream performance are reviewed such as cascade lengths and the impacts of upstream laser types. The paper then discusses some important new S-CDMA capabilities like Increased Power per Code, Active Code Selection and Code Hopping and the advantages they bring to the table.

A capacity analysis will be given that shows that S-CDMA can realistically increase upstream capacity by 25%-50% over an A-TDMA-only system, dependent on plant conditions. Without S-CDMA, it may not be possible to reach 100 Mbps upstream service rates. The findings are reinforced with lessons learned from real-world plant deployments.

Many operators are looking to split nodes in order to increase upstream capacity, even before the roll out of DOCSIS 3.0 services. S-CDMA becomes another tool in our upstream toolkit that we can leverage first for DOCSIS 2.0 systems and that would provide a stepping stone to higher 3.0 rates in the future.

Introduction

Synchronous Code Division Multiple Access, or S-CDMA, is a technology that was added to the DOCSIS 2.0 specification, but up until now, it has not been critically needed or significantly deployed. However, an explosion in home generated content, growth in business services and an extremely competitive environment is putting tremendous pressure on the cable upstream pipe. These and other events and conditions are coming together that will create a Perfect Storm for needed upstream capacity ... and S-CDMA may just be the best near-term answer.

Upstream Bandwidth Needs

Upstream bandwidth needs are being driven by many different factors and trends.

Home Generated Content

Consumer devices have made tremendous progress in recent years with no slow down in sight. Inexpensive digital cameras now offer 10+ megapixel resolution which results in multi-megabit files that typically get transferred upstream as the pictures are shared. As HD capable digital camcorders become inexpensive, they have become mainstream home electronics. They will generate digital video clips that can be shared in either a streaming or bulk transfer mode.

In addition to these consumer devices, the spread of low cost video cameras (a.k.a. webcams) will increase the number of applications like video telephony, video surveillance, "nanny-cams" or "granny-cams" that can be shared over the Internet. Many people will use these webcams for a combination of home security and home monitoring. This can generate a significant amount of sustained streaming traffic on the upstream.

Another important aspect of home generated content is the impact of Social Networking. Our world has become extremely interconnected; driven by technologies like instant messaging, texting on cell phones, and social networking websites like Facebook, MySpace and YouTube. This social networking trend when combined with home generated digital pictures and video content will lead to extensive sharing and hence usage of the upstream.

Another factor that will accelerate these trends is the improving ease of use being introduced into home devices by protocols like Universal Plug and Play (UPnP) and Digital Living Network Alliance (DLNA). It is becoming simple for non-technical people, like Grandma, to connect and share content. This will help create a tremendous amount of user generated content that will be shared over the Internet.

Business Services

A significant growth area for many cable operators is in Business Services. While larger businesses may justify a dedicated Fiber To The Premises (FTTP) connection, many Small and Medium Businesses (SMB) require bandwidths that can be satisfied by DOCSIS 2.0 and 3.0 data rates.

Business application bandwidth needs tend to be very different from residential. Most businesses need symmetric services (i.e. upstream bandwidth equals downstream bandwidth); and Service Level Agreements (SLA) require certain amounts of dedicated throughput. Quite often the operator will keep these SMB customers on separate channels from its residential subscribers.

Business customers tend to be located in clusters, so the impact of additional business services will affect some fiber nodes greatly while others will see little or no impact.

Competitive Environment

For years, cable has had the upper hand over DSL with regards to available bandwidth to the subscriber. Throughout the world, Fiber To The Home (FTTH) or Premises (FTTP) deployments are taking place. Verizon FiOS (Fiber Optic Service) is leading the way in North America with its GPON (Gigabit Passive Optical Network) implementation. As FTTP becomes widely deployed, it will turn up the heat on cable operators.

So how high can today's FTTP deployments eventually go? While the GPON technology enables 1.2 Gbps upstream rates on the access link; current home networking technology will limit residential service tiers to the 100-200 Mbps range. A 1 Gbps service rate will be possible for business and power users that have a direct Gigabit Ethernet connection to the GPON ONT (Optical Network Terminal). There may be other limits in the core network that cap the upstream service rates as well.

All is not lost for the cable operator. The user's perception of the differences between service rates becomes less apparent as data rates increase. The difference between 1 Mbps and 10 Mbps may be significant when compared to the difference between 10 and 100 Mbps services. The speed impact then becomes relatively less significant as you compare 100 Mbps services to a 1 Gbps service. This effect is caused by other bottlenecks in the system (and Internet) becoming prevalent and the access link transfer time becoming insignificant as you move from millisecond to microsecond to nanoseconds. Once you have "enough" bandwidth, then it's the overall experience that will determine consumer satisfaction.

With DOCSIS 3.0 capable of supporting 100+ Mbps upstream service rates, the cable operator will be able to remain competitive with GPON systems, barring some unknown killer application that emerges that actually requires a full 1 Gbps.

4G Wireless Broadband Impacts

Mobile devices have also made tremendous progress in data rates in recent years, which look to continue as the 4G Broadband technologies of WiMax and LTE are deployed. The 4G devices will be rich multimedia devices that both consume and generate pictures and video. Mobile devices today offer 5 Megapixel digital cameras that can stream video clips. This resolution will continue to increase with technological advances.

Some studies have shown that almost 70% of cellular calls are from within the home. We would expect that 4G multimedia usage would be similar. This can have a major impact on a cable operators' plant if much of that traffic is routed thru a home gateway (e.g. via WiFi or 802.11n) or thru a home Femtocell (i.e. a miniaturized, home coverage 4G cell station integrated with your home gateway).

Moore's Law Continues

Over the years, Moore's Law has driven high tech advancements and it looks to continue for the foreseeable future. Simply stated, Moore's Law is the doubling of transistors per device every 18-24 months. This increased technology capacity has resulted in a corresponding improved performance, density and power.

There are other famous high tech "laws" such as Gilder's, Metcalfe's and Nielson's that try to capture the rapid changes and impacts to our networking world. However, it is Moore's Law that is most often quoted and used as a yardstick. In fact, some studies have shown that high speed data broadband service offerings have closely tracked Moore's Law over the last 8-10 years. The evidence above suggests that these trends will continue for the foreseeable future.

Bandwidth Projections

Burst speeds vs. Sustained Rates

Throughout computer history, data networks have been useful because of the nature of statistical multiplexing. By offering a shared resource with high burst rates, users get the impression that they have high data rate services. In reality, data usage is very bursty and average or sustained data rates are significantly lower than the burst rate.

A typical downstream channel might be engineered today to provide 1% concurrency for a 10 Mbps data service. This means that during peak busy hour each active user will get an average of 100 Kbps. Upstream channels today typically operate at a much lower data service rate and a much higher concurrency. The upstream concurrency must factor in the contentions from being a shared pipe as well as a relatively significant amount of VoIP traffic that is non-bursty. So, a 1 Mbps Upstream service rate may use a 10% concurrency resulting in 100 Kbps average usage.

As upstream network speeds increase, the statistical gain also increases. This implies that the concurrency can be reduced as burst speed increases. At 10 Mbps, the upstream concurrency can be reduced to something much closer to downstream concurrency such as 2% which provides an average of 200 Kbps during peak busy hour. As rates increase to 100 Mbps service, this can be supported with a 0.5% concurrency that results in each user getting an average of 500 Kbps during peak busy hour.

The significance of the above is that while the data service rate increased by 100-fold from 1 Mbps to 100 Mbps; the actual bandwidth provided by the operator (i.e. average or sustained rates) only had to increase five-fold from 100 Kbps to 500 Kbps.

Note, this analysis applies only to bursty data applications and does not apply to streaming voice or video. Care must be taken with concurrency if the bulk of upstream traffic does become streaming video.

Not All Subscribers Are Created Equal

Another important aspect of understanding bandwidth needs is to look at the usage across the many different subscribers. Many operators have noted that a relatively few power users consume a proportionately large amount of the bandwidth. These power users may account for anywhere from 1-2% to almost 10% of the subscriber base.

Service Rates – A Five-Year Horizon

As we look at all of the above factors, it seems to indicate that Moore's Law type of growth in consumers' demand for access data rates will continue, and may even accelerate. The implication that demand for access data rates will continue to track Moore's Law means that we'll need to offer 10 times today's service rates in another 5-7 years.

With broadband providers generally offering upstream service rates in the 2-10 Mbps range today, they should be prepared to make 20-100 Mbps services generally available in 5-7 years.

These rates are in line with DOCSIS modem technology. Older DOCSIS 1.1 modems can only support up to a 10 Mbps upstream PHY rate and will presumably be phased out in the 5-7 year time frame. The bulk of DOCSIS devices will probably be 2.0 modems capable of 30 Mbps upstream PHY rates that would support 20-25 Mbps data service rates, which is at the low end of our 5-7 year projections. Power users will use DOCSIS 3.0 technology which can bond up to four upstream channels today to approach, the 100 Mbps service rate. With impairments in the system, it may actually take six or seven bonded upstream channels (which are not available today) to reach and even exceed the 100 Mbps service rate. This will require use of upstream spectrum below 15 MHz which means the use of S-CDMA will be needed. This will be explored further in the paper.

The other piece to the equation that must be taken into account is the required increase in average bandwidth that must accompany higher service rates. As we saw in the previous analysis, this could be on the order of six-fold.

If a typical operator is running a pair of upstream channels today at 10-20 Mbps each, our 5-7 year projection would require at least a half dozen channels providing an aggregate of 150 Mbps. Most cable operators would assume that they'll have to split the upstream nodes to obtain this kind of increase. The remainder of this paper will discuss the S-CDMA technology and how it can help the operator achieve these upstream rates in a single 5-42 MHz spectrum without node splits.

S-CDMA Technology Summary

S-CDMA Overview

Synchronous Code Division Multiple Access, or S-CDMA, is a DOCSIS upstream PHY technology added in DOCSIS 2.0 and enhanced in DOCSIS 3.0. S-CDMA uses orthogonal codes to spread information out during transmission, much as multi-tone modulation uses sub-carriers or even as channel bonding uses completely independent DOCSIS RF carriers to multiplex information. In S-CDMA, each transmitted symbol is stretched out to 128 times the original length, and a 128-bit code is impressed on them such that 128 symbols can be simultaneously transmitted without interfering with each other. Thus S-CDMA has the same inherent capacity in a channel as TDMA.

It is this stretching out of symbols in the time domain that gives S-CDMA its hundredfold advantage over TDMA against impulse noise and short burst noise. [This applies to an event that for example lowers the SNR on a channel by 5 dB or more for up to five or ten (roughly) symbols in TDMA or five to ten chips in S-CDMA.] S-CDMA also adds Trellis Coded Modulation (TCM) to the upstream scheme for additional robustness. Further, a framing structure was added that allows interleaving to be spread over a much greater time frame, which further improves performance against burst noise events (even relatively long burst noise events). Code orthogonality is maintained by synchronizing the transmissions, which can also lead to slight improvements in channel efficiency on small packets via reduction of the guard time and preamble length. The tighter synchronization also contributes to improved performance in impulse and burst noise. A complete description of S-CDMA is found in the DOCSIS RFI specification.

A key benefit of S-CDMA is its robustness against impulse and burst noise, which is a dominant problem in the upstream band below 20 MHz. Another significant benefit of S-CDMA was added in DOCSIS 3.0: the ability to boost the SNR of selected modems on highly attenuated upstream paths by increasing the power per code of individual modems without affecting the operation of the rest of the network. This feature gives S-CDMA an advantage over TDMA for highly attenuated upstream channels such as those encountered in the upper band edge of the upstream where diplex filter roll off is coupled with long amplifier cascades (note there are two diplex filters per amplifier) or the modem is behind multiple splitters such as those found in set-top boxes or across the entire upstream in multiple dwelling units (MDUs) where poor cable and/or many splitters in the path also attenuate the upstream beyond normal levels.

Finally, while not explicitly part of the DOCSIS 2.0 PHY specifications, S-CDMA requires good ingress canceller technology just as do TDMA and A-TDMA. The DOCSIS 2.0 specification provides signal characteristics and hooks to support interference mitigation algorithms that can be implemented by the CMTS. This capability was significantly enhanced in DOCSIS 3.0 systems where a feature called Active Code Selection (ACS) is used by 3.0 modems and CMTS systems with S-CDMA to enhance the ingress cancellation by up to 20 dB. Since these features are critical to the rationale for using S-CDMA, they will be described in more detail in the following subsections.

DOCSIS Considerations

DOCSIS 3.0 vs. DOCSIS 2.0 Operation

While DOCSIS 2.0 S-CDMA provides impulse and burst noise robustness capability, the ingress robustness in 2.0 S-CDMA is limited to cancellation of relatively few and fairly narrowband ingress signals. Further, DOCSIS 2.0 S-CDMA lacks the support of the Increased Power per Code (IPPC) feature. Nonetheless, for upstream channels below 20 MHz with relatively few narrowband ingress signals but replete with impulse and burst noise, 2.0 S-CDMA will be an excellent solution for adding further upstream capacity.

A recommended scenario is to use 2.0 S-CDMA for adding a single additional upstream in the impulse laden portion of the upstream, carefully choosing the particular carrier frequency used to avoid regions where large numbers and/or wideband ingress signals are found. We first evaluated this over a year ago on several live operator fiber nodes. Moreover, Motorola later returned to conduct an extensive comparison for performance below 15 MHz, directly comparing A-TDMA versus S-CDMA. The entire point of that evaluation was to observe the performance of S-CDMA and A-TDMA in the presence of heavy impulse noise and reasonable levels of Ingress Noise. A quick summary shows that DOCSIS 2.0 S-CDMA enabled 32-QAM with lower error rates than A-TDMA achieved using 16-QAM. We could have reduced the Forward Error Correction (FEC) overhead for S-CDMA and gotten additional capacity improvement. We are confident that we'll be able to operate that channel at 64-QAM once tested with the full DOCSIS 3.0 S-CDMA capabilities.

As 3.0 modems are deployed however, S-CDMA channels with full ingress cancellation capability as well as the Increased Power per Code feature can be added near the lower and upper band edges to fill out the upstream spectrum. In this manner, a graceful transition from 2.0 to 3.0 S-CDMA can be effected that uses both existing modems and new upstream spectrum.

Support for DOCSIS 1.1 Modems

In order to mix TDMA and S-CDMA transmissions on the same upstream channel, DOCSIS 2.0 and later versions use separate logical channels for S-CDMA and TDMA transmissions. This adds some overhead to the channel but allows legacy DOCSIS 1.0 and 1.1 modems to share the upstream with S-CDMA.

For moderate impulse noise laden channels, a typical use would be to have DOCSIS 1.x modems operating at, for example, QPSK below 20 MHz, while sharing the channel with 2.0 or higher S-CDMA modems operating at a much higher order QAM. While over time the 1.x modems are likely to disappear from deployment, the use of logical channels permits mixing the oldest and newest modems in a manner to maximize upstream usages for any particular deployment of different types of modems.

RFoG Support for S-CDMA

RF over Glass (RFoG) is a new FTTH technology currently being defined by SCTE that allows existing cable services including DOCSIS to be delivered over the fiber. RFoG is being targeted for green field builds, MDUs and business service applications.

A key issue with RFoG systems is controlling the number of modems transmitting simultaneously in order to prevent the return lasers from being overloaded. S-CDMA support can be accomplished with simple configuration settings in the CMTS. The CMTS need only be configured in the worst case to permit single modems to transmit in each S-CDMA frame. RFoG systems will typically be using the highest upstream rates which results in the shortest transmit times and that helps minimize this issue.

While the need for impulse immunity may be reduced in RFoG systems, the vast majority of impulse noise still comes from the home as we discuss in detail in our Findings from the Field sidebar section at the end of the paper. Since S-CDMA is the best technology to deal with impulse noise, it will still be needed in RFoG systems.

The S-CDMA IPPC feature may be useful in an RFoG system as well. There still may be significant power level variations due to in-home losses and splits and optical link losses as well. IPPC could be helpful in some special situations.

MDU sites fed by RFoG will also benefit from S-CDMA technology. Even though long cascades of diplex filters are gone in RFoG systems, the fiber to an MDU will still be split many times inside the building, often with older coax that is not easily replaced, so the Increased Power per Code feature will still be necessary. Also, the impulse noise in an MDU will be more severe than a single home since the MDU still sums the upstream interference from all dwellings even in the RFoG scenario. So the need for the impulse immunity of S-CDMA remains for MDUs even when connected to an RFoG system.

Getting The Most From S-CDMA

Increased Power per Code

A unique feature of S-CDMA for dealing with highly attenuated upstream channels was added in DOCSIS 3.0: the ability to limit the maximum number of codes used by an individual modem to a lower number than the rest of the modems on the channel, and thereby boosts the power level per code transmitted by the modem.

This 'Increased Power per Code' feature benefits any modem that is pushing the limits of its transmit power. This includes various situations such as modems at the end of long amplifier cascades; DOCSIS Set-top Gateway (DSG) devices that are behind additional splitters in the home; near the edge of the upstream spectrum where diplex filters roll off (e.g. 39-42 MHz); or located in certain MDU environments. This feature allows these borderline modems to use the same higher order QAM at a lower individual throughput as modems on the rest of the network without reducing the capacity of the unimpaired modems while fully utilizing the total capacity of the entire shared channel. A capacity reduction for *all* modems sharing a channel would be required with TDMA operation regardless of sufficient fidelity or transmit power margins available in the majority of modems.

The situation is analogous to multi-tone schemes where a single modem only uses half or a fourth of the available tones, and concentrates its transmitter power on that smaller number of tones to affect an increased SNR per tone, while the remaining tones are used by other modems at the same instant to make normal transmissions.

While this Increased Power per Code feature does reduce the maximum burst capacity of the individual modem using the feature, the limited burst rate is well within typical upstream rates provisioned for cable modems. Further, another DOCSIS 3.0 technology, channel bonding, can be used to augment the maximum burst rate of individual modems using multiple RF channels. Power sharing across the channels will need to be carefully managed, and in extreme cases, modems that need to boost their power per code to the maximum amount may also be limited to a single upstream channel.

In summary, the Increased Power per Code feature overcomes transmit power limitations in highly attenuated return paths and boosts the SNR of affected modems with little to no real capacity degradation without impacting the rest of the modems in the channel. This feature, therefore, also provides an upstream capacity improvement technique that results from using S-CDMA.

Active Code Selection

The reason S-CDMA ingress cancellation performance in DOCSIS 3.0 is superior to that of 2.0 S-CDMA is the active code selection feature of 3.0 S-CDMA. This feature allows the CMTS to select any arbitrary set of individual S-CDMA codes to be used for data transmission. The remaining "unused" codes can be chosen by the CMTS to optimize the response of the upstream receiver for ingress cancellation. This flexible selection of individual codes replaces the fixed selection in the prior version, DOCSIS 2.0. Therefore DOCSIS 3.0 S-CDMA allows choosing the spreading code set for data transmission for optimal ingress cancellation. The figure below shows an example of how the frequency spectra of two S-CDMA codes vary. This is also why the S-CDMA feature of code hopping was added in the original DOCSIS 2.0 specification, as described later.

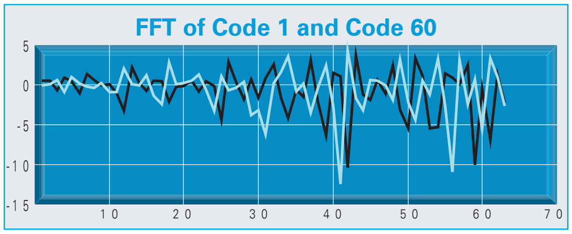
FFT of Code 1 and Code 60

The result is that 3.0 S-CDMA ingress cancellation is equal or better than the excellent ingress cancellation capability of TDMA signal processing that has been demonstrated in the field in modern CMTS receivers. In addition, 3.0 S-CDMA can simultaneously tolerate both heavy impulse/burst noise as well as heavy ingress interference in any portion of the upstream spectrum. Add to this the Increased Power per Code feature of 3.0 S-CDMA, and a more effective total solution for high capacity operation in any portion of the upstream band is available to cable operators using S-CDMA.

Code Hopping

As was seen in the figure above, the frequency spectra of individual codes vary from code to code, and include notches as well as spikes in the spectra. This means that ingress at a particular frequency can affect codes differently by as much as 20 dB. Hence, code hopping was added to the S-CDMA specification to ensure that all modems are fairly treated relative to any particular interference on the plant.

Code hopping is thus an important feature when heavy ingress levels are present. When code hopping is enabled, the codes are systematically shuffled frame to frame so that each modem is given different codes to use. This is done even in the case where a modem's grants fall in the same portion of the frame repetitively. So, code hopping ensures that no single modem is adversely impacted by ingress which could be the case if code assignments remained static.

Real-World Lessons

Findings from the Field

For those that are not faint of heart, we have included an extremely detailed sidebar at the end of the paper describing some of our real-world lessons based on actual data we've obtained from the field. This sidebar takes a closer look at common impairments like micro-reflections, amplitude and group delay phase distortion, ingress noise, common path distortion, impulse noise, and thermal and other random noise sources. The next few sections in our paper will highlight some of these findings.

HFC Impacts: Modulation and Channel Width

A major challenge to implementing wider DOCSIS channel widths, especially 6.4 MHz, is the issue of structural micro-reflections. This impacts DOCSIS 2.0 and DOCSIS 3.0 as well as A-TDMA and S-CDMA systems.

It was long believed that the dominant micro-reflections that DOCSIS had to deal with were created by the customer *home* and little could be done about this type of micro-reflection other than to compensate using an adaptive equalizer. In DOCSIS 2.0 and 3.0, the cable modem transmit equalizer was extended to 24-taps; well beyond the DOCSIS 1.1 CM transmit equalizer of 8-taps. This move alone resulted in tremendously improved equalization performance for DOCSIS 2.0 and 3.0 systems using either a 1.6 MHz or 3.2 MHz channel bandwidth.

However, the extended transmit equalizer has proven to be insufficient in many cases for the 6.4 MHz channel bandwidth. The reason is a combination of the structural micro-reflection and coaxial amplifier cascade depths over N + 5. For the most part the 24-tap DOCSIS 2.0 and 3.0 equalizer can handle either a major structural micro-reflection or a major diplex filter roll off as a result of amplifier cascade greater than five amplifiers deep, but it is the combination of the two major linear impairments that often results in higher than acceptable uncorrected codeword error rate (commonly measured as EQ-MER) for a 64-QAM or in some cases even a 32-QAM modulation.

Ironically all structural micro-reflections are actually caused by the HFC network itself and not the home as previously thought; therefore, it is controllable by the operator. Motorola has been working with CableLabs® to help the operators identify and locate structural micro-reflections. While the operator can technically eliminate the structural micro-reflections, often it is only practical for them to reduce the magnitude of the micro-reflection to the point where it is a non issue from a DOCSIS 2.0 and 3.0 performance perspective. More detail is provided in the Real-World Lessons sidebar.

The distortion due to diplex filter roll-off is an example that can also be dealt with by using a slightly lower carrier frequency. In many cases lowering the carrier frequency less than 1 MHz provides the DOCSIS 2.0 and 3.0 equalizer sufficient dynamic range relief that it can deal sufficiently with the structural micro-reflections.

Impact of Upstream Lasers

A key lesson learned in the field is that operation of the upstream laser is critical. Operation below 15 MHz would be difficult or impossible with certain older FP laser technology. The dynamic range of the return laser can impact both the channel bandwidth and modulation order.

There are basically four to six different types of return lasers in the field today ranging from the most economical FP lasers to the most expensive enhanced DFB lasers. Each advance in laser technology has resulted in improvements in both transmit optical power capabilities and signal-to-noise capabilities. The bottom line is that any given laser has a maximum total input power capability and when that value is exceeded the laser first becomes non-linear. If the input power level increases further over that maximum input power specification then the laser will eventually be clipped.

In the optical world a clipping event means that the laser has been turned off and no further optical energy will be outputted by the laser. The reality is most lasers are seldom clipped but do become non-linear on occasion. These cases of system non-linearity result from a signal with legitimate power combining with ingress noise power. This non-linearity has a much greater effect on the higher order DOCSIS modulations like 32-QAM, 64-QAM and 128-TCM.

Obviously the cable operator can significantly improve the situation by controlling the total ingress noise power present at the input to the laser. Additional improvements in system signal-to-noise performance can also be accomplished by improving the return laser performance by replacing a FP laser, as just one example, with a DFB laser. More information is provided in the Real-World Lessons sidebar.

As operators plan to improve their upstream capacity, they should plan to upgrade the laser technology as appropriate. The DFB laser is the preferred choice today for maximum performance.

The Barren Land – Below 15 MHz

Today, the upstream spectrum below 15 MHz is barren land, typically not being used at all for DOCSIS high-speed return channels. With S-CDMA technology, operators can cultivate this and get a substantial increase in upstream capacity.

Fundamentally the landscape below 15 MHz is fraught with both ingress noise and impulse noise. Poorly shielded coaxial cables are acting as an antenna that easily pick up all available over-the-air transmission signals in the return path frequency range as well as picking up impulse noise that mostly originates from the home.

In order for A-TDMA to be robust in the presence of major impulse noise levels, a narrow 1.6 MHz channel bandwidth is needed. This narrow band makes the TIME SPAN for the FEC markedly longer so it can better protect against the noise. As the DOCSIS channel bandwidth increases above 1.6 MHz, A-TDMA has more susceptibility to impulse noise.

Clearly the Achilles heel for A-TDMA is using a 6.4 MHz channel bandwidth below 15 MHz. A-TDMA is ill prepared to survive in the presence of major impulse noise events (which is far greater below 15 MHz) and where the A-TDMA FEC TIME SPAN is the shortest.

There are other issues with using narrow channels below 15 MHz. It requires at least four of these narrow A-TDMA channels bonded together to achieve the total capacity that a single S-CDMA channel can achieve. Because cable modems are limited today to bonding four upstream channels, this use of multiple narrow channels gets us further away from our goal of trying to reach high upstream rates like 100 Mbps which demands that we use wider channels.

S-CDMA was specifically created to mitigate impulse and burst noise. Since the vast majority of this impulse noise exists below 15 MHz, it is only logical then that S-CDMA be now considered as the technique of choice below 15 MHz.

Capacity Analysis

S-CDMA vs. A-TDMA

So now that we've seen various methods for getting the most from your S-CDMA system, let's quantify these impacts on your available upstream bandwidth capacity. For this analysis, we assume that the entire upstream spectrum is being made available. We also assume that the older DOCSIS 1.1 TDMA technology has been phased out. This analysis focuses on a comparison of alternative DOCSIS 2.0 technologies: S-CDMA vs. A-TDMA.

We know from experience that there is no such thing as a typical cable system. Every system is unique and will have conditions particular to that plant. However, we will give a simple example that shows the potential benefit for rolling out S-CDMA in your system. For this example, we will use a plant with a cascade length of six (i.e. N=6). This is a middle of the road system, one that does not have an excessive cascade length, but is a likely candidate for a node split with the expected bandwidth growth.

In looking at the upstream capacity, we are only considering wider upstream channels of 10 Mbps or higher. We could increase overall upstream capacity with lots of smaller channels (e.g. 1.6MHz QPSK), but this does not get us closer to our goal of offering up to 100 Mbps upstream service rates with existing 3.0 silicon that can only bond up to four channels.

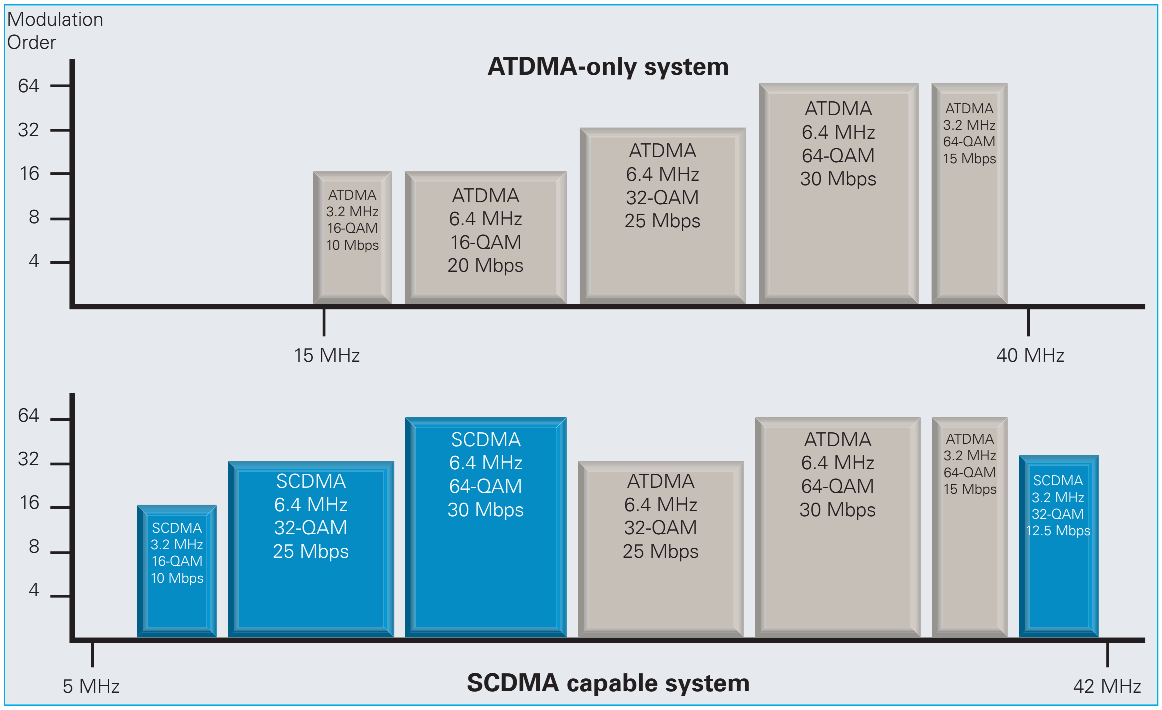

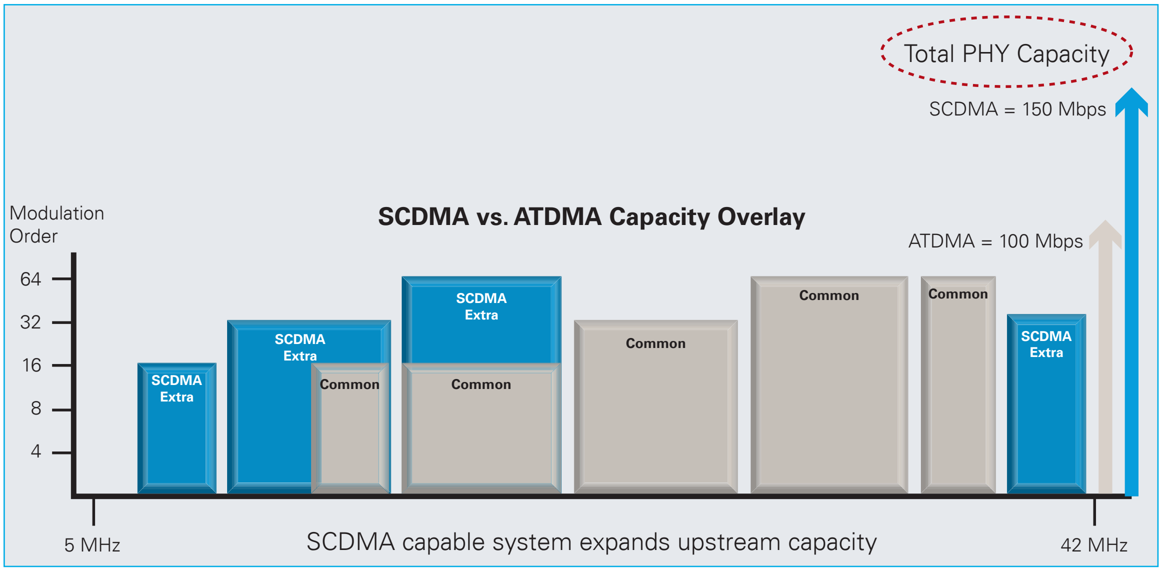

A representation of the channel capacity for an A-TDMA-only system is shown in the first figure above. Using only A-TDMA technology would effectively limit your upstream capacity to the 15-40 MHz spectrum. As we discussed earlier in the paper, ingress noise effectively prevents A-TDMA from operating wider channels under 15 MHz. The noise also reduces the modulation used to 16-QAM for spectrum under 24 MHz. On the upper band edge, it will be tough for A-TDMA to operate in the 39-42 MHz region due to the cumulative group delay from the cascade length.

The representation of the channel capacity of an S-CDMA capable system is the next figure above. As we discussed previously, S-CDMA lets you recover portions of your spectrum that A-TDMA is not able to use. S-CDMA can give you two extra channels below 15 MHz and one additional channel at the high band edge of 42 MHz. In addition to the extra channels, the noise immunity of S-CDMA allows you to use higher order modulation in the 15-24 MHz region for additional upstream capacity. Note that some of this bandwidth gain may be reduced if you are using Active Code Selection and reduce the number of active codes to improve ingress cancellation. For example, using AC=112 instead of the full 128 codes, would reduce that channel's capacity by 12.5%.

We had a chance to compare A-TDMA below 15 MHz to DOCSIS 2.0 S-CDMA on a live MSO cable plant to help us validate the above assumptions. The S-CDMA channel at 32-QAM had significantly lower error rates than A-TDMA at 16-QAM while using identical Forward Error Correction (FEC) values. The S-CDMA could reduce its FEC overhead and get additional capacity increases. The many benefits of DOCSIS 3.0 S-CDMA should allow us to get to 64-QAM and provide even more capacity over A-TDMA. This validates our assumptions above.

The final figure above overlays the A-TDMA-only system on top of the S-CDMA capable system for comparison. The A-TDMA-only system provides a total PHY rate of 100 Mbps across five channels, while S-CDMA enables a PHY rate of almost 150 Mbps over seven channels. Even with a slight reduction in capacity for active code selection, S-CDMA gives about a 40% increase in upstream capacity over an A-TDMA-only system.

But there's more to this story. In the section describing upstream bandwidth demand, we projected that operators will need to support upstream service rates in the 20-100 Mbps range. Note that we showed PHY rates in the picture that must be de-rated by 20-25% to account for DOCSIS overhead and contention in order to obtain usable service data rates (i.e. 10 Mbps PHY rate supports up to 8 Mbps service data rates). With today's silicon supporting four 3.0 bonded channels, the A-TDMA-only system in our example can support service rates to 72 Mbps while the S-CDMA capable system can reach 88 Mbps, a 22% improvement.

Future DOCSIS 3.0 cable modem silicon should be enhanced to allow six to eight bonded upstream channels to allow maximum possible burst rates. This would also imply that an update is needed to the DOCSIS 3.0 specification to require more than four bonded channels. With this next generation silicon, the S-CDMA capable system could reach service rates of ~120 Mbps, a 50% increase over the A-TDMA-only system that maxes out at ~80 Mbps and never reaches our 100 Mbps target rate.

Split Nodes or S-CDMA?

Many operators may presume that they need to split nodes as the next step to increase their upstream bandwidth capacity. With the new nodes, the operator will also have to add additional upstream CMTS receivers as well. With DOCSIS 3.0 bonded services, the operator will likely need four channels on each node after splitting. In reality, most nodes are unbalanced and splitting nodes does not automatically double available bandwidth. If two-thirds of the subscribers are on a single leg, then the effective gain in average bandwidth per subscriber is only 50% for that leg.

As an alternative, the operator could take the additional CMTS receivers needed for the node split and apply them as new S-CDMA channels within the existing upstream as we described in our previous example. This would give an additional 50% upstream capacity, matching the bandwidth per subscriber increase that would be seen from an unbalanced node split. In addition to this, the S-CDMA approach also enables upstream service rates that are significantly higher than if the operator just split nodes and kept A-TDMA-only channels.

S-CDMA has become another tool that we have added to our upstream toolkit. S-CDMA can give the operator up to 50% additional upstream capacity and it enables the operator to offer service rates of 100 Mbps or higher. Node splits still have their place for upstream upgrades. Node splits can double upstream capacity for balanced nodes; it can reduce the cascade length and overall noise levels, resulting in improved upstream capacity and reliability, and it reduces the overall service group size which helps increase the average bandwidth available to each subscriber.

Summary

Upstream bandwidth needs are going to continue to grow. This is driven by factors like home generated content, 4G Broadband wireless and business services. An extremely competitive environment will also add pressure to upstream service rates.

As service rates continue to follow Moore's Law, cable operators will need to generally support 20-100 Mbps services in the next five to seven years. To support these service rates will require six to eight upstream channels with a total PHY capacity of up to 150 Mbps.

To achieve these upstream capacities will require the use of the entire 5-42 MHz spectrum. As we learned from our real-world lessons, there are numerous hurdles to achieving this. A-TDMA is ill equipped to deal with the "barren land" below 15 MHz as well as upper band edge. S-CDMA is truly the modulation technique of choice to deal with the more dominant impulse and burst noise and with moderate levels of ingress noise.

S-CDMA provides other benefits. The 'Increased Power per Code' feature benefits any modem that is pushing the limits of its transmit power, including long amplifier cascades; DSG set-top boxes; upper band edge; and MDU environments. The Active Code Selection capability enables 3.0 S-CDMA to operate robustly in the face of both impulse AND ingress noise.

The upstream capacity analysis of an A-TDMA-only system vs. a S-CDMA capable system was done based on what we learned in the field. In a typical real world environment, S-CDMA can provide up to ~50% additional upstream capacity beyond an A-TDMA only system. Our example shows a S-CDMA PHY rate of almost 150 Mbps vs. 100 Mbps for A-TDMA. S-CDMA also accomplishes this with wider channels, which is critical towards achieving the high service rates that will be needed. In fact, an A-TDMA-only system may max out its service rate around 70-80 Mbps while S-CDMA may enable service rates to 120Mbps.

Many operators have considered node splits as their only option to increase upstream capacity. We have shown that S-CDMA has become an important tool that we have added to our upstream toolkit. S-CDMA can give the operator up to 50% additional upstream capacity and the added S-CDMA channels make it possible to reach the upstream service milestone of 100 Mbps.

In conclusion, S-CDMA is a critical technology needed to increase DOCSIS upstream capacity to remain competitive over the next five to seven years.

Sidebar: Findings From The Field

RF Impairment Definitions

In the end, there are three classes of impairments that all potentially impact a DOCSIS system transmission in the return path:

- **LINEAR IMPAIRMENTS**: Micro-reflections; Amplitude and Group Delay Distortion; Amplitude Slope or Tilt
- **NON-LINEAR IMPAIRMENTS**: Common Path Distortion; Not so Common Path Distortion – Return Laser Clipping
- **TRANSIENT IMPAIRMENTS**: Ingress Noise; Impulse Noise

We'll now take a closer look at each type of impairment:

- **INGRESS NOISE** (CLASSIC) – Narrowband AM carriers such as short-wave radio signals. Ingress interference or noise ranges from **-25 dBc to +15 dBc range.**
- **COMMON PATH DISTORTION** (CPD) – A common occurrence in numerous return path networks is Common Path Distortion (CDP). As the name implies, CPD is the result of a coaxial connector (common to both forward and return paths) becoming corroded and acting as a diode (non-linear device). CPD can easily be identified as the CSO – CTB products will appear in multiples of the forward path channel spacing. CPD products range from **–50 dBc to –20 dBc.** However, what is characteristic of a true laser clipping event is about to occur is that when the laser is within 1 dB of the clipping point, all IMD products as high as **9th** order will suddenly appear and obviously these IMD spurious products are combinations of all return path signals and bears no resemblance at all to common path distortion spurious components.
- **NOT SO COMMON PATH DISTORTION –** This type of distortion is mainly the result of a major ingress noise, impulse noise signal clipping the **Return Laser**. The duration of this type of distortion tends to be brief (less than 100 ms). However, what is characteristic of laser clipping all IMD products as high as **9th** order and they are not simply multiples of the forward path channel spacing.
- **MICRO-REFLECTION** – Micro-reflection is perhaps the most common impairment that exists in every plant by definition. The only difference is whether or not the micro-reflections are significant or not. The most common class of micro-reflections occurs at the lower tap values of the coaxial plant distribution. The lower the tap value the poorer the isolation. Micro-reflections are also by definition frequency dependent and thus not all carrier frequencies and DOCSIS channel bandwidths (symbol rate) are impacted equally. The typical range for micro-reflections is **-45 dBc to -6 dBc.**
- **AMPLITUDE DISTORTION** – There is no questioning the fact that there is amplitude distortion in every plant in the world. Of course there are indeed two major sources of amplitude distortion (i.e. amplitude roll-off) and they are: **COAXIAL CABLE LOSS** – Not a significant impairment in the return path direction; **DIPLEX FILTERS** – The major source of amplitude distortion.
- **GROUP DELAY (PHASE) DISTORTION** – There is no questioning the fact that there is group delay (phase) distortion in every plant in the world. Particularly where there is filtering in a system. In the case of group delay there is really only one major source in a return path plant: **DIPLEX FILTERS.** As was the case with amplitude distortion, the DIPLEX FILTER is the major cause for group delay distortion. In fact the more amplifiers in cascade, the more dramatic the impact of both amplitude and group delay distortion have on a DOCSIS transmission.
- **IMPULSE NOISE** – Impulse noise is also a reality in virtually all return path plants. With that stated there are numerous types of impulse noise that occur on return plant and there are: **POWER LINE (MAINS) IMPULSE NOISE** – This particular impulse noise is often caused by multi-phase electric motors operating. Since they are operating on all three phases, there is the inevitable imbalance in current draw with the result being that energy in multiples of the power line frequency (50 Hz or 60 Hz) are dumped out onto the coaxial plant. In this case the impulse duration can easily be in the **1 ms to 10 ms** range. **Wideband IMPULSE NOISE** – There are multiple sources of this type of impairment. However, the duration of this class of impulse noise or hits, lasts in the **1 ms to 10 ms** range

Life in the Field

"A funny thing happened on the way to the FORUM" is a quote that always fits when having to deal with the real-world implementations for return path transmission systems. In reality the vast majority of nodes in the world are dynamically changing on a regular basis and more often than not the documentation for any specific node is seldom accurate or at the very least outdated.

Given this reality it is always a challenge when going to the field. Just one example of this is attempting to locate a problem structural micro-reflection based on the performance characteristics and then you immediately discover that the network map for the specific node does not reflect what you are actually looking at in the field. Often times we have discovered that the node has already been split and the map doesn't even reflect that fact.

Another harsh reality of the real world is the fact that node splits themselves are an imprecise science to say the least. On one recent trip to the field, we discovered a recently split node that for some reason (having to do with equal homes passed), the split ended up with two nodes with 250 homes passed, but one node was an N + 10 Cascade and the other half of the split was an N + 3 Cascade. Obviously, attempting to locate the structural micro-reflections that often occur as a result of the node splits (segmentations) is problematic at best. The bottom line is that a structural micro-reflection is the direct result of an impedance mismatch in the actual network. These micro-reflections can traverse the coaxial cable as far as 1460 feet, results in a structural micro-reflection time delay of 3500 ns. Simply stated, the DOCSIS 2.0 and 3.0 24-tap equalizer with the MAIN TAP set to TAP-8 has a maximum practical TIME SPAN for the micro-reflection of (N-1)-MT (8) or 15 taps which equals 15 * 195.3125 ns per TAP equals 2929 ns. Even if this micro-reflection had no dispersion and the very last tap of the equalizer could handle the micro-reflection the max time delay for a MR for the DOCSIS 2.0 and 3.0 equalizer is then 16 * 195.3125 ns = 3125 ns. As one can clearly observe the structural micro-reflection in time is simply longer than the DOCSIS 2.0 and 3.0 equalizer can handle. The only solution for this particular case was to minimize the micro-reflection to be better than 35 dB down.

Another harsh reality of the return path is the ingress noise and impulse noise. Typically, up to 90 percent of the total ingress noise power lives below 15 MHz and up to 95 percent of the ingress noise power lives below 20 MHz.

While DOCSIS 2.0 and 3.0 technology possess ingress noise cancellers to deal effectively with the ingress noise interference, the return path laser does not have any immunity to ingress noise power and, therefore, the dynamic range of the laser is typically controlled more by the ingress noise power than any legitimate DOCSIS signal power.

The first line of defense is to attempt to minimize ingress noise power to maintain maximum performance out of the return laser. There are several techniques that are commonly used to achieve this. Typically, since the vast majority of the noise comes in on a poorly shielded coaxial drop cable, the very first line of defense is to replace the older RG-59 coaxial drop cable with RG-6. RG-6 not only has less loss than the older RG-59, but the shielding capability of RG-6 can be as much as 60 dB better than RG-59 when both are brand new, but, more importantly, even 20 year old RG-6 is roughly 20 dB better than brand new RG-59 and 80 dB better than 20 year old RG-59. Obviously many ingress noise problems are effectively dealt with then by the replacement of the drop coaxial cables.

The second approach that is often used comes at the direct expense of return path useable bandwidth. That is to place a 15 MHz high-pass filter directly at the input to the return laser. Since the vast majority of the ingress noise power lives below 15 MHz, this technique does dramatically improve the return path laser dynamic range at the direct expense of bandwidth, which is now reduced to the 15 MHz to 42 MHz range. As we strive to maximize our upstream bandwidth over the long haul and leverage the 5-15 MHz spectrum, this approach is not a long-term solution.

The preferred solution is to replace the return laser with an improved technology laser. The chart below displays the performance differences given the various technology lasers available. The total dynamic range capabilities of five different laser technologies are shown. One can observe that the older and cheaper FP lasers have such a small operation window of performance acceptance that it is displayed here in this paper only to stress the point that the FP laser cannot provide any meaningful bandwidth for modern DOCSIS 2.0 and 3.0 services.

It is also apparent that while the isolated FP laser and the enhanced isolated FP lasers are significant improvements over the original FP laser technology, in the long run the more appropriate selection moving forward, especially for a DOCSIS 3.0 system, ought to be either a DFB laser or a high-power DFB laser.

Note, in order to contain our discussion, we will not cover digital return systems in this paper. Digital return systems are equal in performance to DFB

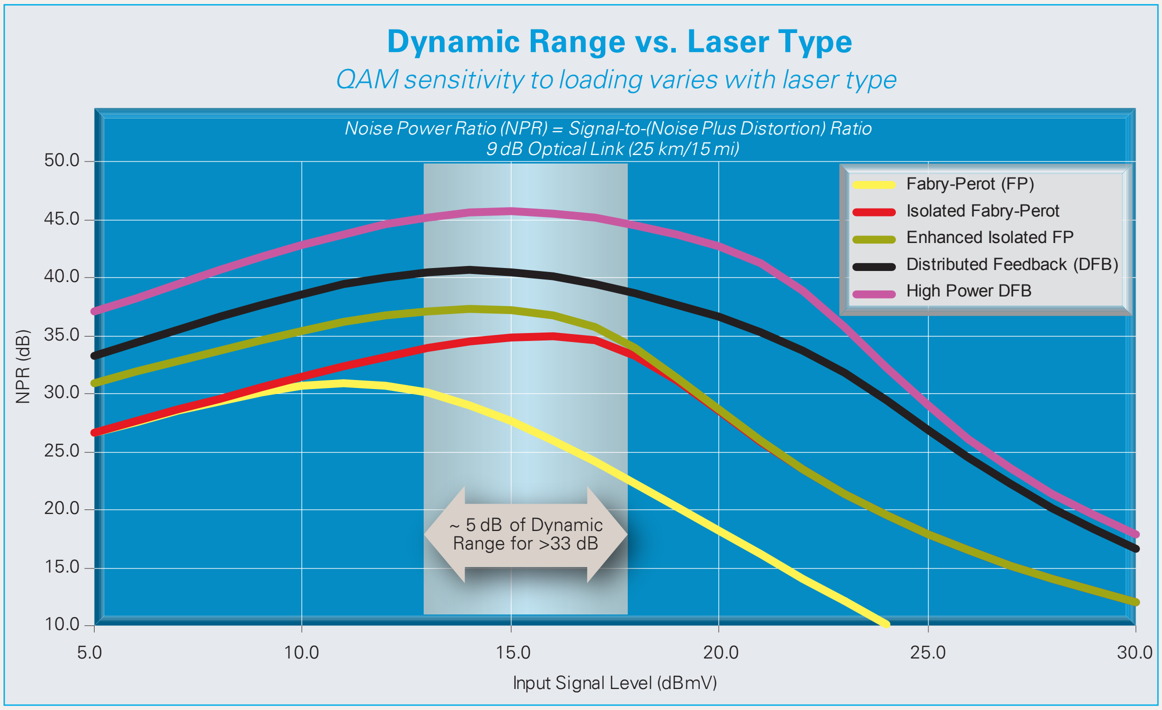

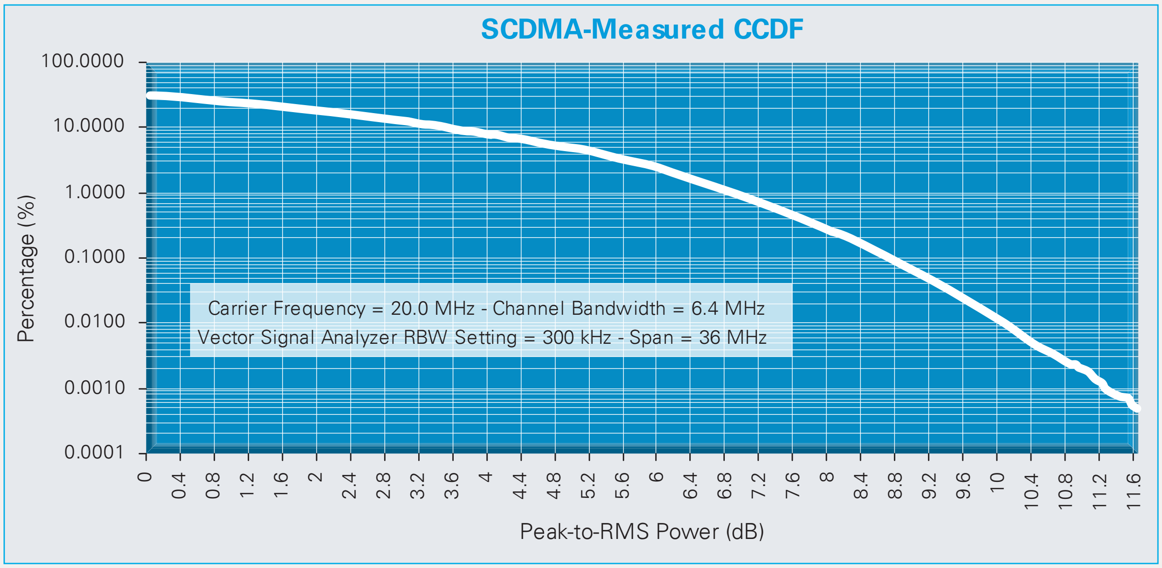

lasers. However, because there is no standard for digital return systems as of this writing, their use has been limited.

The final harsh reality of real-world lessons learned in the field is that DOCSIS 2.0 and DOCSIS 3.0 S-CDMA modulation technology has a Peak-to-RMS power factor that rivals an Additive White Gaussian Noise distribution. Simply stated a single S-CDMA signal stresses the dynamic range of the laser about as much as the total ingress noise power from a peak perspective anyway. Therefore, multiple S-CDMA channels really require a minimum of an isolated FP laser and preferably a DFB laser to truly take advantage of the capabilities in performance improvements that S-CDMA affords the cable operator.

The other figure above is the Complimentary Cumulative Density Function for a live S-CDMA burst measured in the Motorola Marlborough HFC lab given a Agilent 89640A Vector Signal Analyzer with a RBW = 300 kHz and a Span = 36 MHz of a S-CDMA 20 MHz carrier using a 6.4 MHz channel bandwidth signal.

More on the Barren Land – Below 15MHz

The CATV industry has all but surrendered using any bandwidth below 15MHz. Fundamentally the landscape below 15 MHz is fraught with both ingress noise and impulse noise. The mechanism that couples the vast majority of both ingress noise and impulse noise into the CATV return path network is poor shielding of an older coaxial drop cable. This poorly shielded coaxial cable is acting as an antenna that easily picks up all available over-the-air transmission signals in the return path frequency range as well as picking up impulse noise that mostly originates from the home.

The energy, whether it be ingress noise from over-the-air signal transmission or impulse noise emanating from the home, enters the CATV network on the ground and, generally speaking, this type of coupling mechanism tends to act as a low-pass filter function as the effective coupling decreases beyond 13 MHz. In fact, both the impulse noise and ingress noise levels drop off by roughly 18 dB per octave. Therefore, by 26 MHz the interference power of both is down by 18 dB and by 39 MHz the interference signals are down by 36 dB.

Given that the interference factor drops off significantly beyond 15 MHz or so, it is easy to understand why the CATV industry has simply avoided using DOCSIS signal transmission below 15 MHz up to now. However, the operators are now in need of more bandwidth and the bandwidth that has yet to be exploited is the 5 MHz to 15 MHz spectrum.

With the creation of DOCSIS 2.0 the concept of ingress noise cancellation became an accepted function for all DOCSIS 2.0 CMTS receivers upstream. Ironically, though, one will not find a single requirement for ingress noise cancellation in either the DOCSIS 2.0 or DOCSIS 3.0 standards. Moreover, with the development of DOCSIS 2.0 came the introduction of S-CDMA and it was generally understood from the very beginning that S-CDMA was adopted because of its ability to handle impulse noise. Since the vast majority of impulse noise exists below 15 MHz, it is only logical that S-CDMA be now considered as the technique of choice below 15 MHz.

The A-TDMA Dilemma

Generally speaking TDMA/A-TDMA has been the modulation technique of choice up until now, but it has been primarily used above 15 or 20 MHz where the power of the impulse noise is low enough that the A-DTMA FEC capabilities could easily handle impulse noise. However, below 15 MHz the impulse noise is far more powerful and a far greater stress on the existing A-TDMA FEC capabilities.

While it is generally accepted CATV industry wide that DOCSIS 2.0 A-TDMA can handle virtually any ingress noise interference, it is also understood that the same robustness can’t be assumed for impulse noise. In fact, in order for A-TDMA to be extremely robust in the presence of major impulse noise levels, one needs to select a 1.6 MHz channel bandwidth which makes the time span for the FEC markedly longer for such a narrow bandwidth.

With each increase in DOCSIS channel bandwidth above 1.6 MHz A-TDMA becomes more susceptible to impulse noise. In fact, a characteristic failing of an A-TDMA modulation technique in the presence of a major impulse noise event is a significant percentage of FEC uncorrected codewords with a very small percentage of FEC corrected codewords as the duration of the impulse noise events is a greater time that the A-TDMA FEC can span.

Clearly the Achilles heel for A-TDMA using a 6.4 MHz channel bandwidth below 15 MHz is that it’s ill prepared to survive in the presence of major impulse noise events in which the power of the impulse noise is far greater when the A-TDMA FEC time span is the shortest.

In recent years network maintenance programs have done a reasonable job of managing ingress noise to manageable levels from a total power perspective. The return lasers are no longer being over-driven at the rate they were just five years ago. However, the same cannot be said for impulse noise as it is typically created in the home. As stated above the max power for most impulse noise events are in the 5 MHz to 13 MHz frequency region of the return path.

With A-TDMA operating at a 3.2 MHz channel bandwidth the practical maximum data signaling rate would optimistically be 32-QAM at 2560 ksym/s (12.8 Mbps). Realistically speaking it is more likely that the upper limit with acceptable error rates would be 16-QAM at 2560 ksym/s (10.24 Mbps) or 32-QAM at 1280 ksym/s (7.68 Mbps).

Even if you add four channels of 1.6 MHz, the best A-TDMA could offer for a data signaling rate would be F1 at 5.12 Mbps + F2 at 6.4 Mbps + F3 at 6.4 Mbps + F4 at 7.68 Mbps for a total bandwidth of 25.6 Mbps. This would require a DOCSIS 3.0 upstream channel bonding to accomplish what a S-CDMA can do using a single 6.4 MHz channel bandwidth to achieve a 32-QAM or 25.6 Mbps. With known DOCSIS 3.0 cable modem implementations supporting just four bonded upstream channels, using four narrow A-TDMA channels also has the limitation that its maximum PHY burst rate is 25.6 Mbps. With S-CDMA, its wideband channels can be bonded with other wideband channels to obtain PHY burst rates in excess of 100 Mbps.

In conclusion, S-CDMA is truly the modulation technique of choice to deal with the more dominant impulse noise and with moderate levels of ingress noise. To support growing bandwidth needs with wideband channels (e.g. 20.48 Mbps or higher rate) below 15 MHz, the CATV industry really has no choice but to opt for S-CDMA.

Abbreviations and Acronyms

802.11 – The IEEE standard associated with home or premise wireless networking

802.11n – most current version of 802.11 technologies

4G – Fourth Generation broadband wireless, e.g. WiMax and LTE

ACS – Active Code Selection

AGC – Automatic Gain Control

A-TDMA – Advanced Time Division Multiple Access

BER – Bit Error Rate

CATV – Community Antenna Television, early name for cable industry

CCN – Composite Carrier to Noise

CM – DOCSIS Cable Modem

CMTS – DOCSIS Cable Modem Termination Systems in Head End

CPD – Common Path Distortion

CSO – Composite Second-Order distortion level

CTB – Composite Third-Order distortion level

DFB – Distributed Feedback laser

DLNA – Digital Living Network Alliance

DOCSIS – Data over Cable Service Interface Specification

DOCSIS 1.1, 2.0, 3.0 – various versions of the DOCSIS RFI specification

DSL – Generic reference to all Digital Subscriber Line technologies, e.g. ADSL, VDSL

FEC – Forward Error Correction

FemtoCELL – a small, home coverage cell station integrated with a home gateway

FiOS – Fiber Optic Service, Verizon's name for its FTTP service

FP – Fabry-Perot laser

FTTH – Fiber-to-the-Home

FTTP – Fiber-to-the-Premises

GPON – Gigabit Passive Optical Network

HFC – Hybrid Fiber Coaxial system

HD – High Definition video content

IPPC – Increased Power per Code

MDU – Multiple Dwelling Units

MPEG – Moving Picture Experts Group

MPEG-4 – Next generation MPEG digital video compressions standard

N+6 – HFC architecture with Node plus six actives

OLT – Optimal Line Termination (GPON termination equipment in the Central Office)

ONT – Optimal Network Termination (GPON termination equipment at the home)

PHY – Physical layer

PON – Passive Optical Network

QAM – Quadrature Amplitude Modulation

QPSK – Quadrature Phase Shift Key modulation

RF – Radio Frequency

RFI – DOCSIS Radio Frequency Interface specifications

RFoG – RF over Glass

S-CDMA – Synchronous Code Division Multiple Access

SD – Standard Definition video content

SDV – Switched Digital Video

SLA – Service Level Agreement

SMB – Small and Medium Businesses

SNR – Signal-to-Noise Ratio

TCM – Trellis Coded Modulation

TDMA – Time Division Multiple Access

UPnP – Universal Plug and Play

VOD – Video On-Demand

VoIP – Voice over Internet Protocols

WiFi – Industry alliance organization promoting 802.11 standardization

Acknowledgements

The authors would like to thank the following people for their inputs and technical review: Rob Howald, Dean Stoneback and Mike Cookish of Motorola; Victor Hou, Bruce Currivan, Rich Prodan, Jonathan Min, Tom Kolze and Ernie Bahm of Broadcom.

This paper was originally presented at the SCTE Conference on Emerging Technologies – April 2009 – Washington, D.C.

MOTOROLA

Motorola Inc. www.motorola.com

09/09