

DOCSIS™ 2.0

Technical Analysis of DOCSIS 2.0

white paper

Abstract

The purpose of CableLabs® first Data Over Cable System Interface Specifications (DOCSIS) - DOCSIS 1.0 and DOCSIS 1.1, were to respectively enable residential data services and voice services over a single Internet Protocol (IP) cable infrastructure. The 1.0 specification defined the upstream and downstream physical and data link layers necessary to transmit over shared multiple-access cable IP networks. DOCSIS 1.0 specified the basic Quality of Service (QoS) features required to offer tiered services based on rate-limits, and was later enhanced to support minimum guaranteed rates. The DOCSIS 1.1 specification introduced support for constant bit rate services, which greatly enhanced the QoS feature set, and somewhat improved the robustness of the return path, which allowed twice the bandwidth, while providing full backward compatibility with the 1.0 specification.

In December 2001, CableLabs® released the first version of the DOCSIS 2.0 specification. The primary objective of DOCSIS 2.0 is to enhance upstream spectral efficiency, which requires additional robustness. This paper's objective is to investigate the new features introduced in DOCSIS 2.0, by closely examining its benefits to legacy 1.0 and 1.1 cable modems (CMs), and 2.0 CMs.

DOCSIS 2.0 achieves the goal of increasing upstream spectral efficiency and robustness by enhancing the 1.x (DOCSIS 1.0 and 1.1) time division multiple access (TDMA) modulation encoding method, renaming it Advanced-TDMA (A-TDMA), and by introducing a new upstream modulation encoding method known as synchronized code division multiple access (S-CDMA). DOCSIS 2.0 requires that CMs and cable modem termination systems (CMTSs) support both A-TDMA and S-CDMA, thereby leaving the choice of enabling either or both methods on the DOCSIS channels to the operator. This paper analyzes both encoding schemes to help the reader better understand the how they should be enabled in the network.

The Evolution to DOCSIS 2.0

Introduction

This paper presents facts about the benefits of A-TDMA and S-CDMA (as embodied in the DOCSIS 2.0 RFI), and the relative advantages of one vs. the other. It was written to help multiple system operators (MSOs) better understand these technologies so that decisions may be made about whether and how to deploy them.

Document Overview

This paper first describes some of the improvements that an A-TDMA product will have, compared to the existing DOCSIS 1.1 product. All of the improvements described apply to legacy DOCSIS 1.X cable modems. Many of these improvements go beyond DOCSIS 2.0 requirements.

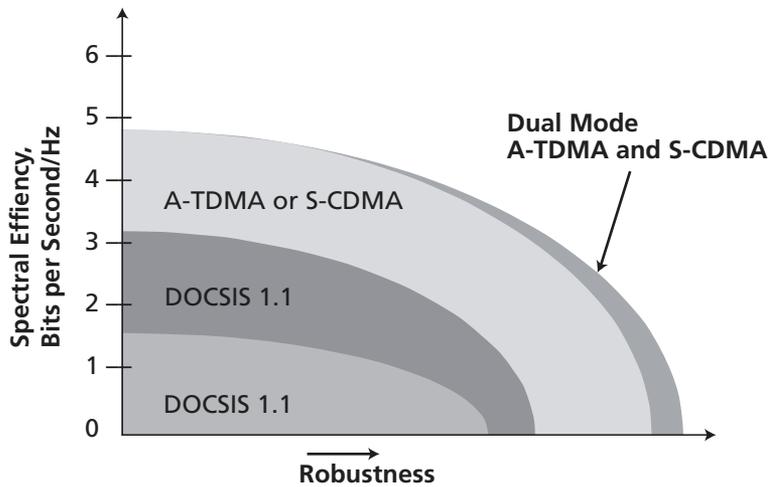


Figure 1: Evolution of DOCSIS robustness & spectral efficiency

- Ingress Cancellation
- Improved Receive Equalization
- Improved Burst Acquisition
- Impulse Noise Mitigation (improved forward error correction, or FEC)

The paper then goes on to compare the two modulations (A-TDMA and S-CDMA) in the following areas:

- Dynamic Range and Timing Sensitivity
- Impulse Noise and Ingress
- Backward Compatibility and Interoperability of A-TDMA and S-CDMA
- Scheduling Efficiencies

Lastly, a conclusion is presented as to how an operator may use these facts in evolving towards DOCSIS 2.0.

Figure 1 shows a graphical illustration of the evolution of robustness and spectral efficiency. Note that spectral efficiency is quantifiable in terms of bits per second per Hertz. Robustness is essential for *allowing* operation at higher bits per second but cannot be quantified without a fully defined channel model and is therefore more subjective. The intent of the illustration is to show that there is a large step up in robustness by evolving to S-CDMA or A-TDMA and a smaller incremental benefit results from switching between modes according to plant conditions.



The Evolution of Advanced PHY

Work on Advanced PHY techniques began in March of 1998 under the auspices of IEEE 802.14a and culminated with the release of the DOCSIS 2.0 specification by CableLabs®. The fundamental improvements that were desired by MSOs from an advanced standard were: 1) increased capacity, 2) increased robustness to radio frequency (RF) upstream impairments, and 3) no degradation to existing DOCSIS deployed systems. Increased robustness was desired both for improving the robustness of existing networks, but also to make possible the higher orders of modulation and symbol rates of Advanced PHY, since the larger constellations and rates require either cleaner RF channels or greater robustness to be reliably deployed. Some of the robustness enhancements of Advanced PHY are laid out in the specification, such as increased FEC, while other robustness enhancements are implemented in the receivers such as ingress cancellation, and thus are proprietary in nature, but also apply to existing DOCSIS networks. Hence, these improvements can also be used to increase the reliability of medium bandwidth channels, such as the 16-QAM, 3.2MHz channels, available in existing legacy 1.X DOCSIS modems, as long as an Advanced PHY CMTS is deployed which supports the new robustness features.

On the other hand, the capacity increase provided by the larger constellations and rates requires that *both* the CMTS and the CM have Advanced PHY features. It should be noted that even partial Advanced PHY deployments improve the entire network capacity since the Advanced PHY modems are using less resources per modem for the same provisioned level of service. Lastly, the requirement for Advanced PHY to not degrade existing services means that the new technology must integrate seamlessly with existing networks and not create additional overhead or other bandwidth consumption that reduces network capacity.

A-TDMA 'Single Ended' Features

While many of the features of DOCSIS 2.0 (such as S-CDMA) do not provide any improvements to legacy modems, there are four robustness improvements that are 'single ended', i.e. may be obtained only by upgrading the CMTS and apply to all DOCSIS cable modems. These provide benefits long before DOCSIS 2.0 cable modems will be ubiquitous simply by upgrading the CMTS.

- Ingress Cancellation Filter
- Improved Receive Equalization
- Improved Burst Acquisition
- Improved Error Correction for Impulses

These improvements will not only improve reliability and capacity in existing upstream channels, but will also make new RF spectrum available to existing DOCSIS 1.x modems, thereby greatly expanding the upstream capacity available to *existing DOCSIS systems*. Each robustness improvement is detailed below.

Ingress Cancellation Filter

The most significant improvement in robustness, ingress cancellation, is actually not part of the DOCSIS 2.0 specification, but should be found in some form in all Advanced PHY CMTSs to support the higher order modulations. This is a digital filter that adaptively responds to narrow band and wide band ingress or common path distortion (CPD) and filters it out (reference Figure 2).

The ingress cancellation filter (ICF) determines the nature of the upstream ingress by analyzing the channel in-between packet bursts. Cancellation coefficients are computed by a digital signal processor (DSP). These coefficients are updated up to 200 times per second to handle time-varying ingress. The ICF is integrated with the receive equalizer to optimize overall channel response.

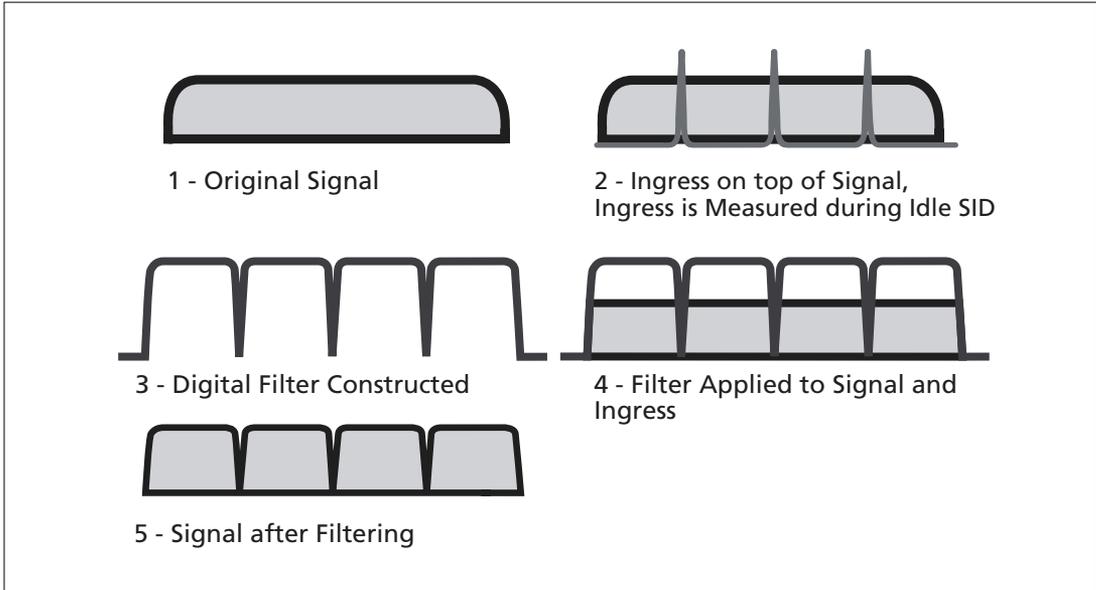


Figure 2: The Process of the ICF filtering out three carrier wave (CW) tones

The ICF is able to cancel or mitigate the effects of narrow band and wide band interference as well as CPD (a form of wide band interference due to non-linear components in the plant). An example of this is shown in Figure 3. An example of multiple CW ingress that can be effectively cancelled by a commercially available A-TDMA burst receiver is shown in Figure 4 along with the 16-QAM constellation after filtering (Figure 5). Table 1 below contains test results from a variety of ingress types and levels.

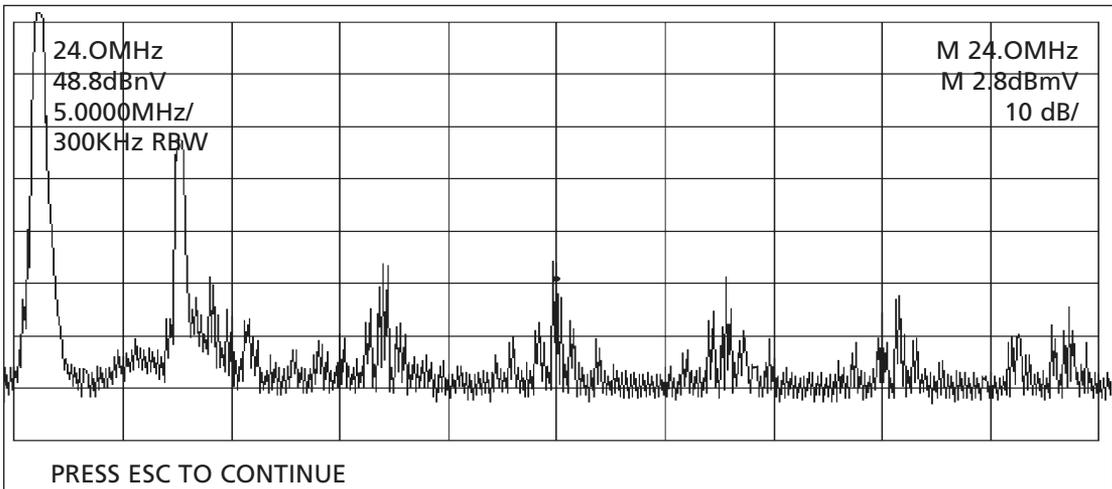


Figure 3: CPD Ingress Example

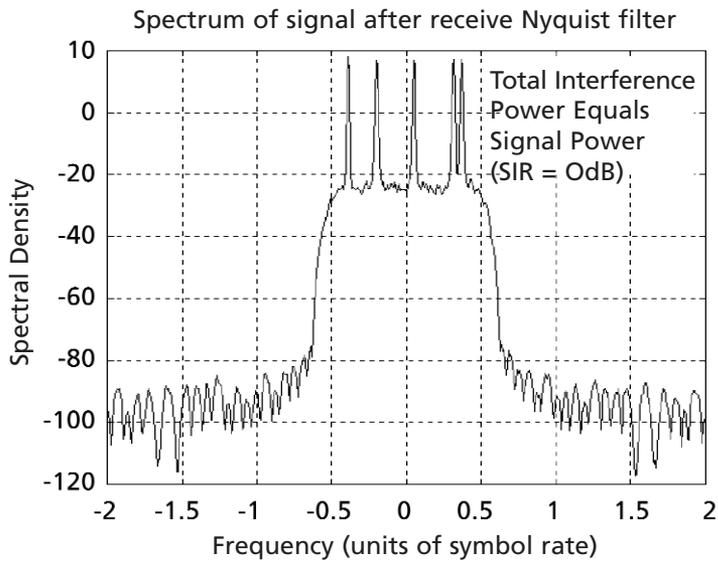


Figure 4: Signal Spectrum with 5 ingressors

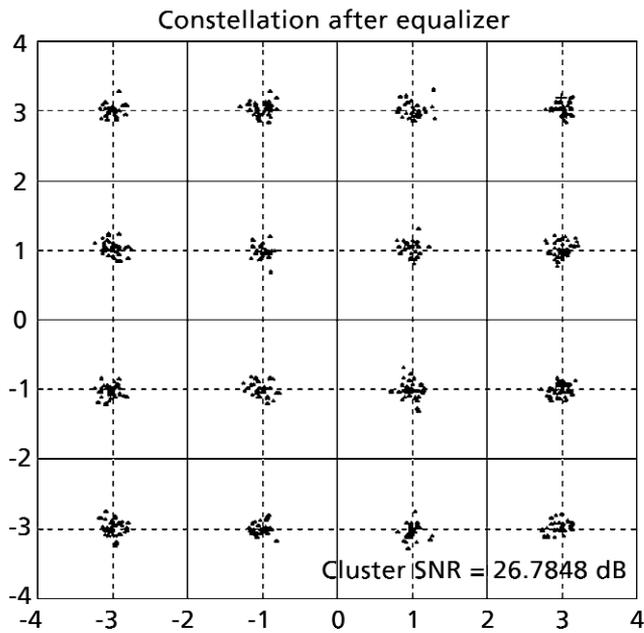


Figure 5: Constellation after ingress removal

Impairment Type	Required CIR without ICF	Required CIR with ICF
Narrow band ingress	20dB	-10dB
5 narrow band ingressors	20dB	0dB
2 wide band ingressors	20dB	9dB
CPD	20dB	5dB

Table 1: Carrier to ingress ratio (CIR) required for 16-QAM, with and without ICF

Improved Receive Equalization

Equalization is used to mitigate the effects of frequency dependent attenuation, delay and multipath on both the upstream and downstream in DOCSIS. In the downstream, equalization is performed at the cable modem receiver. Since there is a continuous downstream signal that is coming from a single source (the CMTS) the equalization can be accomplished completely within the cable modem receiver.

DOCSIS 1.1 and 2.0 implement *upstream* equalization using pre-distortion rather than relying on *receive* equalization. The reason for this is that a large number of preamble symbols are necessary in order to train the equalizer in receive-only equalization¹. Since each CM burst will experience different upstream distortions, if each data burst was preceded by a preamble long enough to train the equalizer, then much of the upstream bandwidth would be wasted, especially for small packets.

Pre-distortion Equalization

On the upstream, the CMTS sees sequential bursts coming from potentially thousands of cable modems. The hybrid fiber coax (HFC) plant (see Figure 6) distorts each burst differently since the bursts travel through different paths and plant elements. To equalize the burst, the receiver must use different equalizer coefficients. It is impractical to store all of these coefficients at the CMTS and load them burst-by-burst into the equalizer. Instead the coefficients are sent to the CMs so that they may *pre-distort* the upstream bursts (see Figure 7). Once they arrive at the CMTS, the bursts will be undistorted, as the HFC plant will reverse the effect of the pre-distortion. Therefore the primary purpose of the CMTS receiver equalizer is to *measure* the distortion and calculate the necessary coefficients for the cable modem based on ranging bursts. In addition, the receive equalizer is still active at the CMTS on data bursts and can help mitigate transient channel distortion.

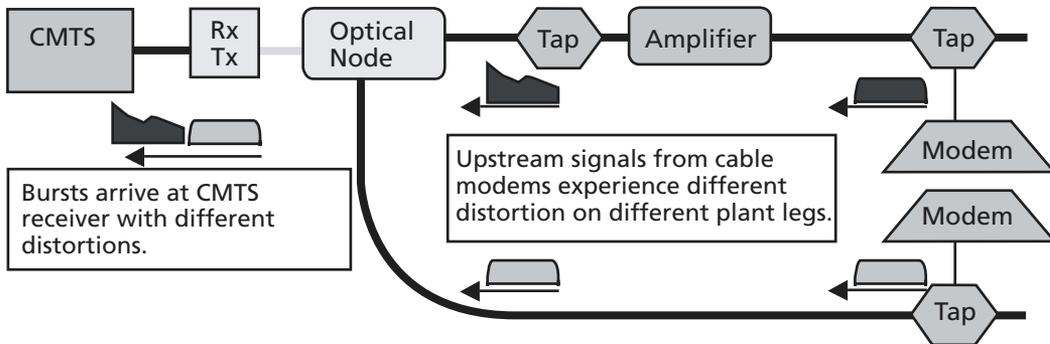


Figure 6: No Pre-Distortion. Bursts must be equalized at the receiver by using a training pattern of adequate length

1. In the downstream equalization is accomplished during initial synchronization and then changes slowly thereafter. Therefore the overhead for equalizer training is insignificant.

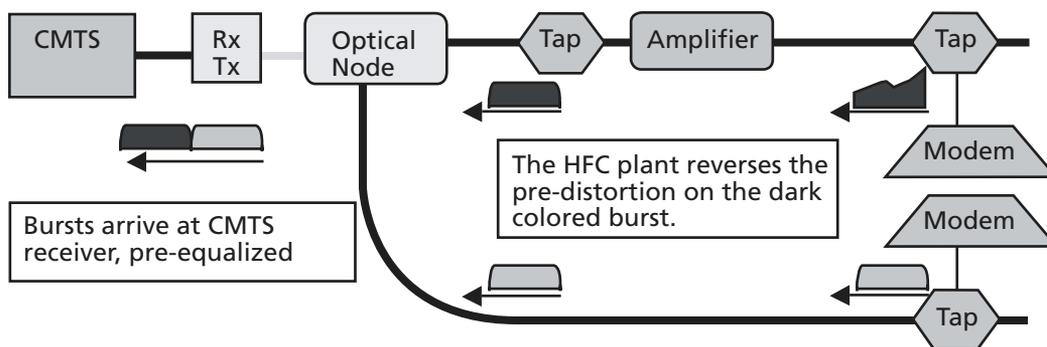


Figure 7: Pre-Distortion used. Bursts are pre-distorted at the cable modem transmitter and the plant “equalizes” them prior to the CMTS receiver.

“24-Tap” Equalization

In DOCSIS 1.1, the pre-distortion is defined for 8 taps. The A-TDMA CMTS has a 24 tap receive equalizer. It would seem that without an A-TDMA CM with 24 tap pre-distortion, that the receiver 24 tap equalizer is useless. In fact, 24 tap receive equalization may be used on a burst by burst basis, due to the improvements that have been made on the A-TDMA CMTS burst acquisition (see section on burst acquisition below). Although maximum pre-equalization is enabled when both the CMTS and CM have a matching number of taps, a higher order receive equalizer will enhance performance in a single ended fashion.

Consequently, an A-TDMA receive equalizer has the capability of compensating for multipath that is *4 times the duration²* of the multipath that can be handled by DOCSIS 1.1 CMTSs (see Figure 8).

Improved Burst Acquisition

The new A-TDMA burst receiver has a greatly improved burst acquisition capability. This was necessitated by higher order constellations, which requires an increased precision in the estimation of acquisition parameters. Increased precision acquisition may be accomplished by long preambles at the expense of efficiency. Instead the A-TDMA burst receiver has a new robust method of acquisition, which is accomplished in a minimum number of symbols. This applies to low order modulation, such as QPSK and 16-QAM, allowing acquisition in the presence of impulse noise and shorter preambles. In addition, the equalizer will train on the entire preamble, which allows the receiver-only equalization described in the previous section (see Table 2 for a summary).

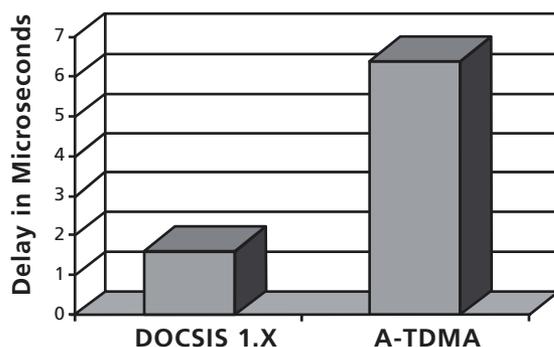


Figure 8: Enhanced ability to tolerate long delay multipath.

2. The burst receiver has the main tap offset from the center at tap 8, which allows 16 trailing taps vs. 4 trailing taps with the standard burst receiver, i.e. 4 times the number of taps.



STANDARD PHY BURST ACQUISITION

- Standard burst receivers acquired signal parameters off of sections of the preamble in *series*. If any section of the preamble is hit by impulses, *acquisition fails*.
- Equalizer training was done *after* the preamble. If multipath is bad, *acquisition fails*.
- In short, acquisition was the weakest link—not the data forward error correction.

ADVANCED PHY BURST ACQUISITION

- Carrier and timing lock, power estimates, equalizer training and constellation phase lock are all done simultaneously. This allows shorter preambles (20 symbols) and/or robust acquisition with impairments.
- Reduction in the implementation loss to a fraction of a dB from theoretical, which means that legacy modems will be able to operate in higher additive white Gaussian noise (AWGN) noise levels than previously possible. Depending on the constellation size, symbol rate, and packet error rate being measured, up to 2.3dB improvement in AWGN performance is available with new A-TDMA burst receiver.

Table 2: Standard vs. Advanced Burst Acquisition

Improved Forward Error Correction for Impulse Noise

DOCSIS 1.X will allow the correction of 10 errored bytes per Reed Solomon (RS) block (T=10). DOCSIS 2.0 allows correction of 16 bytes per Reed Solomon block (T=16). To obtain T=16 performance requires that a DOCSIS 2.0 cable modem is used in conjunction with a DOCSIS 2.0 CMTS.

However, the A-TDMA receiver has a new capability that offers a single ended improvement to FEC. This capability is called *erasure* correction. Erasure correction is possible when the location of errors within the Reed Solomon block is known. [1] Using erasure correction, up to 20 bytes with errors may be corrected for DOCSIS 1.X modems (up to twice as much correction power). Importantly, performance gains due to erasure correction *do not require Advanced PHY cable modems*. The technique works most effectively with impulse or burst noise where the location of errors can be inferred from detection of the impulse event at the demodulator.

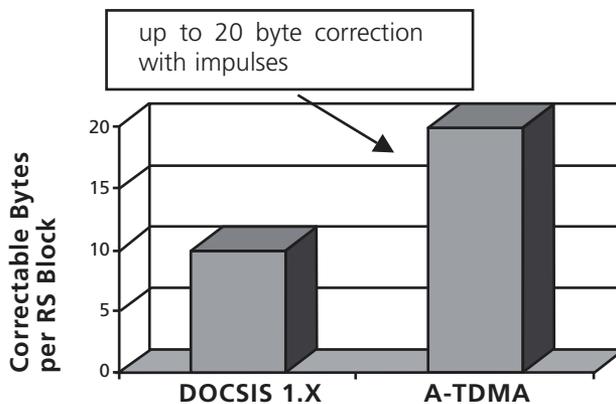


Figure 9: Erasure correction improves burst noise performance



Technical Comparison between S-CDMA and A-TDMA

When examining modulation technologies, it is important to realize that there is not any technology that can exceed theoretical maximums such as the Nyquist and Shannon limits. In addition, if the technologies have been architected and implemented well, there is a tendency for the performance of each approach to converge towards these theoretical limits. In particular, it has been shown that if all other system parameters and coding are equal, then all modulation technologies will have identical performance in AWGN.

Early History of the Advanced PHY Standard

Some perspective may be gained by a short discussion of the early history of the Advanced PHY standard that culminated in DOCSIS 2.0.

As stated earlier, in March of 1998, IEEE 802.14a began work on improving the DOCSIS upstream modulation for increased robustness and bandwidth. There was early agreement that the DOCSIS downstream³ performance was adequate and should remain untouched. At that time, the major battle was between selection of the S-CDMA proposal by Terayon and the variable constellation orthogonal frequency division multiplexing (VCOFDM) proposed by Ultracom (A-TDMA was assumed to be included). The battle hinged, in part, on the disadvantage of S-CDMA requiring tight timing requirements vs. the disadvantage of VCOFDM requiring dynamic constellation adjustments. Other factors, such as Terayon having a deployed HFC S-CDMA system vs. Ultracom's mostly theoretical proposal tipped the scales in favor of S-CDMA.

Technologically agnostic members of the committee demonstrated that all three proposals could provide roughly equivalent performance with each having advantages under specific operating conditions [2]. The same equivalence of performance existed between the Advanced TDMA proposal (by Broadcom and Texas Instruments) when compared to the modulation approaches of S-CDMA and VCOFDM. The main advantage of A-TDMA was that it was an incremental enhancement over the existing DOCSIS TDMA approach. Despite the similarity of performance between A-TDMA and S-CDMA, there are some advantages and disadvantages that result, partly because of the tradeoffs made in the specifications details.

S-CDMA Advantages

The following section assumes some knowledge of how S-CDMA works as specified in the DOCSIS 2.0 RFI. The RFI specification may be found at: <http://www.cablemodem.com/specifications.html>.

Impulse Noise

The long duration symbols of S-CDMA provide robustness in the presence of certain types of impulse noise. Long duration symbols lead to a requirement that multiple symbols be simultaneously sent to maintain aggregate bandwidth. As will be seen, this fact leads to a *crossover point* where S-CDMA outperforms TDMA on one side of this crossover and TDMA outperforms S-CDMA on the other side. The underlying mechanism of the effect is illustrated graphically in Figure 10 and Figure 11. In a sense, the S-CDMA code space is "overloaded" with a high power impulse and transmission must move to a lower order constellation (e.g. quadrature phase shift keying or QPSK), which will reduce throughput. A-TDMA, on the other hand, will experience errors at lower impulse power. This crossover is shown in Figure 12.

3. Based on ITU-T J.83, Digital multi-programme systems for television, sound and data services for cable distribution.

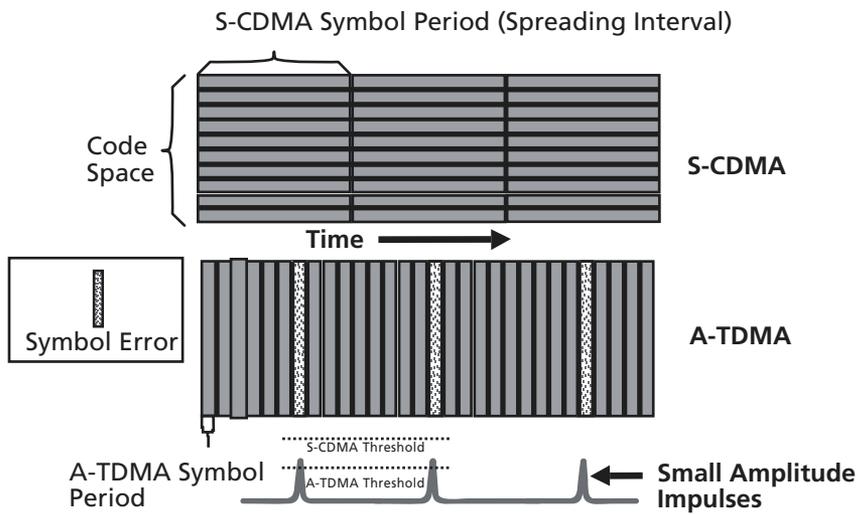


Figure 10: S-CDMA is robust vs. short duration and medium level impulses

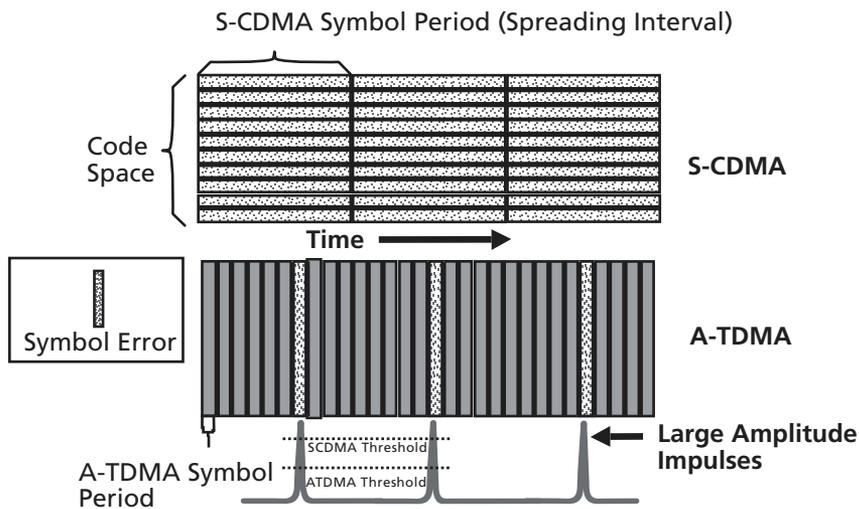
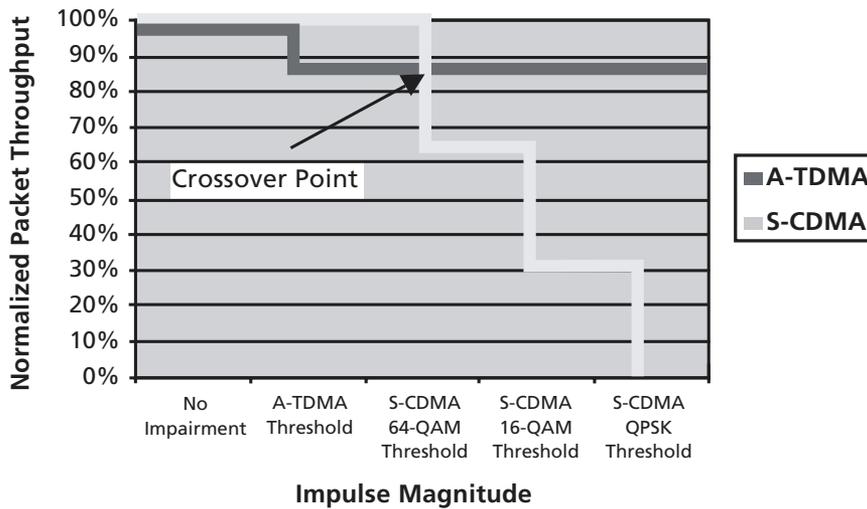


Figure 11: TDMA is robust against high power impulses of short duration

S-CDMA excels in short duration (around 1us) medium amplitude, high repetition rate impulses, and with small packets. S-CDMA tolerates these impulses (as shown in the Figure 10) due to the long duration of the S-CDMA symbols that spread the short impulse energy over the whole symbol⁴. On the other hand, spreading does not help if the impulse energy is high enough to corrupt the symbol (as shown in the Figure 11), despite the spreading effect. In general, most or all of the symbols that are simultaneously sent will be simultaneously corrupted if impulses are high in power. At that point, TDMA has an advantage because the number of corrupted symbols is limited as only one symbol per unit time.

4. As the impulse approaches the symbol duration this spreading advantage declines. Once the impulse is as long as the symbol period there is no spreading.



*Parameters adjusted to maintain <0.1% packet error

Figure 12: Crossover of A-TDMA and S-CDMA impulse noise rejection

Long Impulses

S-CDMA also has an advantage against long duration (greater than 5-10 μ s), large impulses and with small packets. For long impulses, an important factor in S-CDMA impulse noise performance is the spreading interval factor 'K' and the number of codes per minislot (CPMS). If K is large and the CPMS are small, the number of codes simultaneously used for any burst may be small. Another way of looking at this is that the burst is stretched out as far in time as possible, making the impulse duration a small fraction of the burst duration. Since a large impulse will destroy almost all codes that are simultaneously sent, then it is best to send as few codes simultaneously as possible (2-4 codes per minislot). The few symbols that are corrupted can easily be corrected by FEC. As an example⁵, see Figure 13. Assuming that the channel is operating in 64-QAM at 2.56Mps, if the S-CDMA mode is at K=32 (the maximum allowed) spreading intervals per frame and the number of codes per minislot is two, a 64 byte packet may be transmitted in two minislots (excluding preamble). The duration of the frame is 0.39 μ s/chip x 128 x 32 = 1.6ms. The maximum allowable duration of a burst with the maximum Reed Solomon correction of T=16 is 250 us (5 spreading intervals). The sensitivity of S-CDMA burst performance is dependent on the codes per minislot and the K factor (reference graphic of S-CDMA burst performance vs. K in Figure 15).

On the other hand, with A-TDMA, a small packet is sent in a short amount of time (see Figure 14). For a channel operating in 64-QAM at 2.56Mps, a 64 byte packet is sent (excluding preamble) in 33 μ s. With the maximum Reed Solomon correction of T=16, the impulse will corrupt 16 bytes and break the error correction if it is 8.3 μ s in length or greater. Therefore, for the case with S-CDMA optimized for maximum tolerance to burst noise, S-CDMA can handle bursts that are 30 (250/8.3) times longer than A-TDMA. However, there is an impact to dynamic range as a result of operation in the most burst tolerant S-CDMA mode, as will be seen below.

5. In this example the MAC overhead and preamble is ignored. Interleaving is off since it won't improve performance.

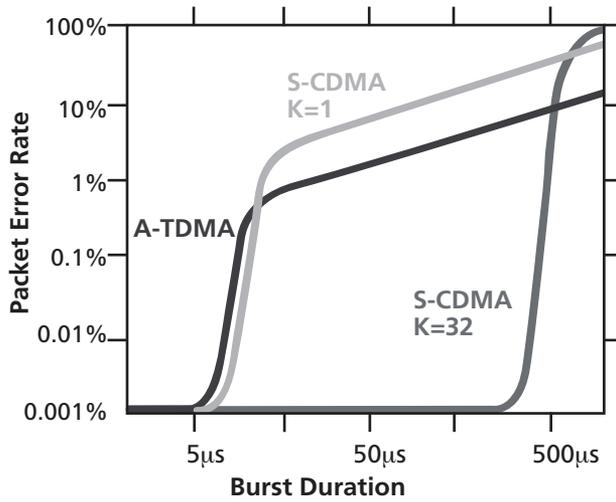


Figure 15: Long Burst Performance of Small Packets (S-CDMA K=32 vs. S-CDMA K=1 vs. A-TDMA)

A-TDMA Advantages

Short Duration High Amplitude Impulses.

Short duration high amplitude impulses are handled better by A-TDMA, as once the impulse is large enough, it will corrupt all S-CDMA codes in a spreading interval (Figure 11). At a high enough repetition rate, error correction will not be able to compensate for these errors, regardless of interleaving. A-TDMA can handle 10 to 100 times the repetition rate since only a single symbol is sent per unit time, causing only one symbol to be corrupted per unit time.

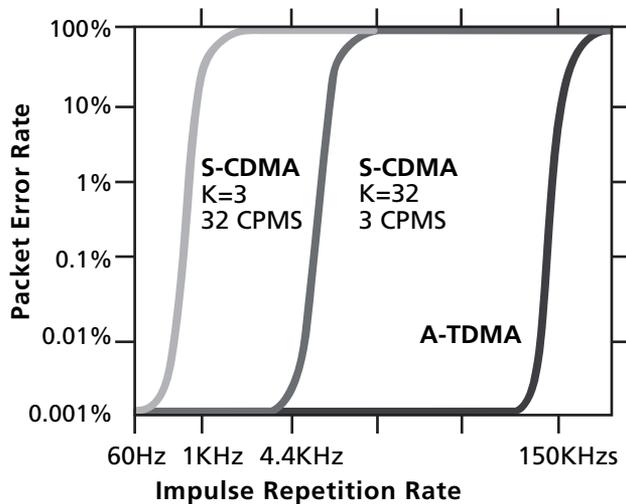


Figure 16: High amplitude impulses are handled effectively by A-TDMA.



The way this works is that after decorrelation at the S-CDMA demodulator the impulse energy is spread over all 128 CDMA 'chips' which make up a single CDMA 'symbol', so that the impulse energy is 1/128th as large. This reduced energy degrades the signal-to-noise ratio (SNR) of each of the 128 simultaneously demodulated orthogonal sequences that make up all the symbols. This works well at mitigating the impulse until the amplitude becomes so great that even 1/128th of the impulse degrades the symbol SNR at the demodulator such that the signal cannot be recovered. All 128 sequences, or 'symbols', suffer the same fate and one large impulse destroys them all.

As an example, consider the same channel as in the previous example, 64-QAM at 2.56Msps and T=16. The packet size will not be critical in this case. S-CDMA will be set to maximum impulse noise tolerance settings of 2 codes per minislot and K=32. Assuming 1us impulses at a level high enough to corrupt both A-TDMA and S-CDMA symbols, at low repetition rates both S-CDMA and A-TDMA FEC will be able to correct the errors. As the impulse rate increases, the S-CDMA system will experience errors when 7 impulses occur within one frame of 1.6ms, i.e. at 4.4KHz. Alternatively, A-TDMA will start to experience errors at 150KHz (Figure 16).

TDMA dynamic range

S-CDMA Low Power Limitation

As we have seen, operation in S-CDMA mode has better impulse noise handling if the number of codes per minislot is low and the K factor is high. The lowest number of codes per minislot allowed is 2. Therefore, this is also the best for impulse noise immunity. However the DOCSIS 2.0 specification requires that the dynamic range of the S-CDMA modem be from 8dBmV to 53dBmV *independent of modulation order and the number of codes per minislot*. This is a critical specification that results in a low power limit that varies according to the number of codes per minislot. The DOCSIS 2.0 specification defines the minimum power with all codes active as:

$$\text{Minimum Upstream Power}_{2\text{cpms}} = 8\text{dBmV} + 10 \cdot \log_{10}(128/2) = 26\text{dBmV}$$

At the minimum power limit, this results in a reduction of dynamic range by 18dB.

S-CDMA High Power Case

Due to the high peak to average nature of S-CDMA signals, there is a required power backoff from maximum power. Note this is similar to the power backoff that was needed for 16-QAM and higher order QAM in DOCSIS 1.X. However S-CDMA requires the same power backoff for all QAM modes including QPSK (also called 4-QAM). Therefore, a S-CDMA modem that is operating in QPSK must operate at a maximum power level of 53dBmV. A-TDMA modem operating in QPSK may operate at 58dBmV.

S-CDMA Dynamic Range vs. TDMA Dynamic Range (see dynamic range graph next page)

The analysis below will assume the case of 2 codes per minislot for S-CDMA. This is the most robust case for impulse noise, but the most limited in dynamic range. Increased dynamic range may be had at the expense of S-CDMA's advantage in impulse noise resistance. This trade-off will be examined later.

There is a detailed analysis of the required dynamic range in the return path (upstream) in the book, "Broadband Return Systems for HFC CATV Networks" by Donald Raskin and Dean Stoneback. The range is extremely large (49dB) until some techniques are applied to reduce the range. One of the techniques is feeder equalization, which is not universally applied to HFC systems. If feeder equalization is applied (among other techniques), then the required dynamic range is 34dB. Most of the remaining variance is found in the customer in-house wiring and splitters, which is not easily subject to control by the MSO. It is worthwhile to compare this range requirement to the actual dynamic range of cable modems in DOCSIS 2.0:



- The total dynamic range of TDMA modems is 58dBmV-8dBmV = **50dB**.
- The total dynamic range for S-CDMA modems is 53dBmV-26dBmV = **27dB**.

Note that in DOCSIS 1.X, the dynamic range was made large to accommodate a wide range of channel bandwidths, from 2,560 ksym/sec to 160 ksym/sec. The reason channel bandwidth matters is that if *uniform spectral density* is desired in the upstream and if multiple bandwidth channels are to coexist, the low bandwidth channels must operate at lower powers. For each doubling of bandwidth, an effective loss of 3dB is experienced for the *system* dynamic range. Over the channel bandwidth range of DOCSIS 1.X, there are four doublings in bandwidth, resulting in $3 \times 4 = 12\text{dB}$ effective reduction in system bandwidth. Since A-TDMA allows operation at 5,120 ksym/sec, there are 5 bandwidth doublings, equating to a 15dB reduction. Therefore the *effective* dynamic range (vs. total dynamic range) of TDMA is $50\text{dB} - 15\text{dB} = 35\text{dB}$, which is 1dB greater than the Raskin and Stoneback recommendations.

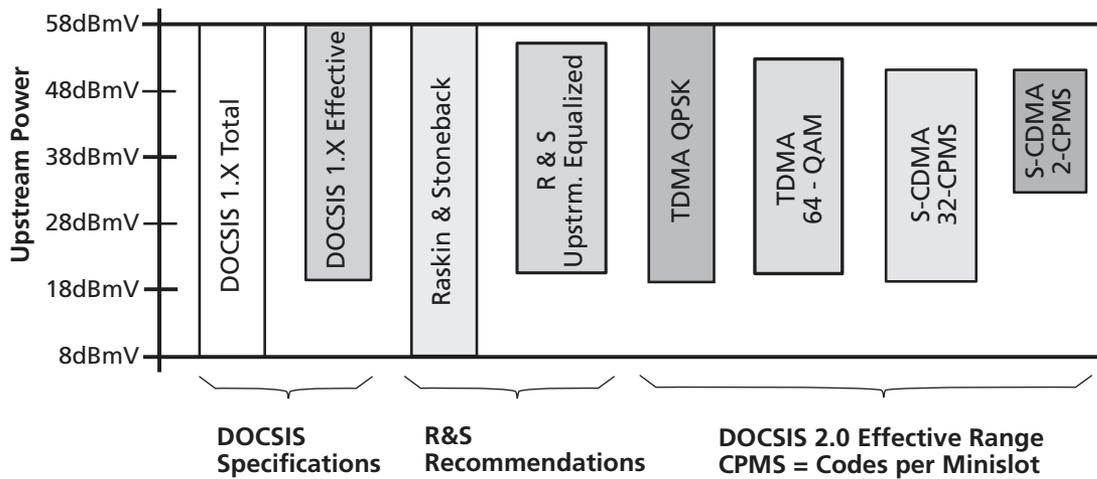


Figure 17: Dynamic Range Implications of S-CDMA

It was understood during the creation of the DOCSIS 2.0 specification, that S-CDMA dynamic range is limited. Therefore operation below 1,280 ksym/sec is prohibited when in S-CDMA mode. The argument is that the Advanced PHY benefits are needed only for wider band channels. Therefore the *effective* dynamic range of S-CDMA is reduced from the *total* dynamic range by only 6dB, since only two doublings of bandwidth are allowed. The *effective* dynamic range of S-CDMA is $27\text{dB} - 6\text{dB} = 21\text{dB}$.

Increasing the number of codes per minislot above 2 will increase the dynamic range on the low power end of S-CDMA modems. With each doubling of codes per minislot, the low power range is increased by 3dB. However, this improvement is obtained at the expense of a tradeoff in robustness against impulse noise. Because the maximum number of codes is 32, which is 4 doublings, the low-end dynamic range is improved by 12dB. The result of using the maximum codes per minislot is an effective dynamic range of $21\text{dB} + 12\text{dB} = 33\text{dB}$.



S-CDMA Differential Code Power and Intercode Interference

With S-CDMA it is important that all modems transmit at the power required by the CMTS. Differential code powers will cause degradation in the codes with low power due to minor timing differences in the codes with high power (relative to the CMTS). The S-CDMA timing budget assumes equal power in all the codes. Therefore if the dynamic range limitations cause different modems to have unequal code power at the CMTS, the timing induced non-orthogonality will cause the lower power codes to experience intercode interference from the higher power codes.

Effect of Dynamic Range in a Cable Plant

A cable modem should transmit at a high enough power as required by the CMTS, otherwise the signal will end up closer to the noise floor. In this case a single modem may force the entire plant to switch to a lower QAM level in order for the 'challenged' modem to operate. In the worst case the CM will be out of the CMTS receive range. At the other extreme, the CM is unable to turn down its transmitter to the CMTS receive level. This will cause too much RF power directed to the upstream laser and may also cause clipping, depending on the laser operating margin. In the extreme, the power will be outside of the CMTS receive range.

Operators should evaluate their plant characteristics and determine the optimal operation as a tradeoff between length and amplitude of impulse noise and dynamic range requirements of the cable plant.

One proposed solution is to switch all modems operating in S-CDMA mode to A-TDMA mode when the dynamic range is needed. This requires the CMTS to obtain the cable modem RF power levels, monitor these levels and switch individual modems to the alternate logical channel for A-TDMA when needed.

Quantization Noise Funneling

Noise funneling in S-CDMA has been addressed by the DOCSIS 2.0 specification, but it is worth noting. S-CDMA allows a variation in the number of simultaneously transmitting modems from 1 to 64. (TDMA only allows a single modem transmitting at one time). As a result, in the 64 simultaneously transmitting case, there is additional upstream noise caused by the large number of modems transmitting 2 codes apiece. The most fundamental source of noise is due to the quantization noise added by the digital to analog converter (DAC). This noise is created by the fact that a DAC has a finite number of bits of resolution. The least significant bit (LSB) will cause a change in the output signal that is a step change with an error compared to the 'ideal' signal. This noise is usually identical to AWGN and can be treated as such. In DOCSIS 2.0 the requirements on the DAC were effectively tightened by 2 to 3 bits of resolution to handle this effect. Nonetheless, operating in A-TDMA mode or in S-CDMA mode with a high number of codes per minislots will lower the funneling noise and decrease the implementation loss⁸.

Statistical Multiplexing Advantages

It is possible to operate all modems, from 1.0 to 2.0 in TDMA mode. It is not possible to operate 1.X modems in S-CDMA mode. Therefore, if any 1.x legacy modems coexist on the same channel with modems operating in S-CDMA mode, the channel must operate in 'dual' mode. This requires a minimum of two 'logical'⁹ channels. Losses in statistical multiplexing efficiency are experienced when a channel is split into sub-channels. This effect may be reduced but not eliminated by intelligent scheduling. In addition, dual Map of minislots (MAP) and upstream channel descriptor (UCD) sets must be sent and upstream contention regions must be segregated.

8. Quantization noise effects combined with limited dynamic range may make low numbers of minislots in S-CDMA a less desirable mode.

9. Logical channels were created in DOCSIS 2.0 to allow coexistence between different modulation modes.



TDMA Relaxed Timing Requirements (reference Appendix B)

S-CDMA requires accurate timing due to the upstream and downstream synchronization requirements. Timing variations can be caused by temperature shifts and possible wind loading effects in the plant. S-CDMA requires no more than a 2ns timing error. Station ranging must therefore be used to adjust timing before a 2ns error can build up. Appendix B discusses the timing changes that may exist in HFC systems and how they may impact the frequency of station maintenance in S-CDMA.

Sources of HFC Impairments

It is instructive to understand the sources of the various upstream impairments described above, and also to understand how A-TDMA and S-CDMA will handle them [3].

Impulse Noise (time varying noise):

- Long Duration (~10ms) Large Impulses (0dBc) - One source of this type of noise is Impulse Pay-per-View Polling. Older upstream signaling from set-top boxes may be unbalanced. These devices were not designed with automatic gain control (AGC) and were usually set at high levels to ensure the signals were above other noise sources. This high signal level causes cross-compression of the upstream lasers and causes high magnitude impulse noise over all frequencies simultaneously (with lower frequencies degraded more than higher frequencies).
- Medium Duration (~100 μ s) - May be caused by load switching using mechanical contacts. Another source is a loose F-connector making intermittent contact, usually due to wind loading.
- Short Duration Impulses (<1 μ s) at 60Hz or Harmonic - These are caused by power line related sources, such as arcing on transformers and thyristors. Motors and appliances such as hair dryers can also generate these impairments.

S-CDMA will spread very short impulses and perform better than TDMA unless the impulses are large in magnitude, then TDMA will perform better. For long bursts, S-CDMA (in general) will spread packets out further in time, thereby making FEC more effective. For long bursts (temporal dispersal) S-CDMA must use low numbers of codes per minislot at a concurrent loss of dynamic range.

Ingress (frequency varying noise):

- Citizens band (CB), Ham Radio and Spurs - These are relatively narrow in bandwidth (<20kHz). These spurs are usually caused by leakage from local oscillators, etc. The CB and Ham Radio are fixed in frequency but keyed on and off over time.
- CPD and Impulse Noise with Harmonic content - This noise is wider in bandwidth but still has repetitive peaks in the frequency domain. Although impulse noise should be in the time varying noise category, some impulse noise has a harmonic signature in the frequency domain. This noise can benefit from frequency domain filtering.

A-TDMA has well known and relatively easy to implement frequency domain filtering. Specifically, one filter used in an A-TDMA design will improve performance from narrow band interference by 30dB. S-CDMA can theoretically have similar performance but the implementation is more complex [2]. The MSO will be well advised to examine this non-specified but critical feature in selection of a CMTS solution.



Summary of S-CDMA and A-TDMA Comparison

S-CDMA and A-TDMA have very similar performance in terms of robustness and bandwidth. The most significant differences between the two are:

- S-CDMA may be configured to have better impulse noise resistance for most impulse types, at the expense of dynamic range.
- S-CDMA has shorter preamble requirements, providing an advantage in small packet conditions.
- S-CDMA has tight timing requirements, which may require more frequent station ranging.
- A-TDMA has a better dynamic range under all modulation settings.
- A-TDMA is backwards compatible with all legacy modems, which results in better statistical multiplexing.
- A-TDMA has implementation advantages in ingress cancellation performance.

Technical Analysis of DOCSIS 2.0 - Conclusion

DOCSIS 2.0 clearly adds a number of enhancements to improve the robustness and capacity of the upstream. *It is important to note that DOCSIS 2.0 only provides the tools for obtaining these benefits.* Given this, a crucial component to improving robustness and bandwidth is the software that utilizes DOCSIS 2.0 features to intelligently adapt to plant conditions. Finely tuned ingress filter DSP software, look-ahead channel hopping, ingress categorization and adaptation, and optimized fallback algorithms are just some of the requirements of intelligent PHY control.

Assuming intelligent PHY software[4], both A-TDMA and S-CDMA excel under somewhat different ingress environments. As was seen in section 2 of this paper, an A-TDMA CMTS provides a variety of robustness improvements to deployed DOCSIS 1.x modems.

The major advantages of DOCSIS 2.0 may be obtained with CMTS that comply with *either* the A-TDMA *or* the S-CDMA requirements. If a cable operator has exclusively deployed DOCSIS 2.0 cable modems and the HFC plant dynamic range complies with the S-CDMA requirements, it may well be desirable to operate in S-CDMA mode or A-TDMA mode depending on the plant conditions. On the other hand, if the HFC plant has a mixture of 1.x and 2.0 cable modems, it may be desirable that an MSO operate in A-TDMA mode.

Ultimately the choice of a CMTS should be based on a wide variety of considerations:

- Intelligent PHY Control Software
- Redundancy and Reliability
- Capacity and Throughput
- Integration with Voice Services
- Integration with Provisioning and Network Management Applications
- Advanced PHY Capability
- Carrier-Class Edge Routing Capability
- Advanced QoS Capability



Appendix A - Acronyms

ACG	Automatic Gain Control
A-TDMA	Advanced Time Division Multiple Access
ATP	Acceptance Test Procedure
AWGN	Additive White Gaussian Noise
CATV	Community Access Television
CB	Citizens Band
CIR	Carrier to Ingress Ratio
CM	Cable Modem
CMTS	Cable Modem Terminating System
CPD	Common Path Distortion (non-linear mixing products)
CPMS	Codes Per Mini-Slot
CTE	Coefficient of Thermal Expansion
DAC	Digital to Analog Converter
DOCSIS	Data Over Cable System Interface Specification
DSP	Digital Signal Processor
FEC	Forward Error Correction
HFC	Hybrid Fiber-Coax
IEEE	Institute of Electrical and Electronics Engineers
ITU	International Telecommunication Union
LSB	Least Significant Bit
MAC	Media Access Control
MAP	Map of minislots (abbreviation)
MSO	Multiple System Operator
OPL	Optical Path Length
PICS	Protocol Implementation Conformance Statement
PHY	Physical layer (abbreviation)
QAM	Quadrature Amplitude Multiplexing
QPSK	Quadrature Phase Shift Keying
RF	Radio Frequency
RS	Reed Solomon (type of forward error correction)
S-CDMA	Synchronized Code Division Multiple Access
SNR	Signal-to-Noise Ratio
TEP	Test Execution Procedure
UCD	Upstream Channel Descriptor
VCOFDM	Variable Constellation Orthogonal Frequency Division Multiplexing



Appendix B - Calculations on HFC Optical Path Length Changes and The Effect on Station Maintenance for S-CDMA in DOCSIS 2.0

Optical Path Length Change with Temperature

Optical Path Length Change Calculations

In an S-CDMA system, it is critical that the upstream codes from the CM be precisely timed so that they arrive at the CMTS with other codes from other CMs. The CMs must add delay such that all codes are aligned to within +/-2ns. Due to changes in the OPL of the HFC system over temperature, periodic ranging is required to keep the packets aligned. Periodic maintenance must be done depending on how quickly the optical path length (OPL) (or delay) changes with time.

The OPL is related to fiber index (n) and length (l) as: $OPL = 1n$

According to Corning Glass Works, the change in refractive index of fiber over temperature is approximately the same as $\Delta n/T$ of fused silica (optical fiber is fused silica doped with Germanium). In addition the change in physical length of fiber over temperature is approximately the same as the $\Delta l/T$ of fused silica.

The effective refractive index of fiber is: $n = 1.47$

The change in refractive index of fused silica over temperature is,

$$\Delta n/T = 1.28 \times 10^{-5}/^{\circ}\text{C} \text{ (from Corning)}$$

$$\Delta n/T = 1 \times 10^{-5}/^{\circ}\text{C} @ 589\text{nm} \text{ (from Oriel Instruments)}$$

The Coefficient of Thermal Expansion of fused silica is:

$$\text{CTE} = \Delta l/l - T = 5.5 \times 10^{-7}/^{\circ}\text{C}$$

The optical path length change will be:

$$\Delta OPL = \Delta n * n + n * \Delta l$$

Normalizing for length and delta temperature results in:

$$\Delta OPL/l_T = (\Delta l/l - T) * n + \Delta n = \text{CTE} * n + \Delta n$$

$$= 5.5 \times 10^{-7}/^{\circ}\text{C} * 1.47 + 1.28 \times 10^{-5}/^{\circ}\text{C}$$

$$= 1.36 \times 10^{-5}/^{\circ}\text{C} \text{ or}$$

13.6 millimeters per kilometer degree centigrade using Corning's $\Delta n/T$. Using Oriel's $\Delta n/T$ we obtain 10.6 mm/km or 10.6 ppm/ $^{\circ}\text{C}$ (parts per million per degree C). It can be seen that the majority of the change is due to the refractive index change and not the physical change of the fiber length. Given the extraordinary low coefficient of thermal expansion (CTE) of fused silica, this is not surprising. (It turns out that the Oriel number is closer to reality therefore, this will be used.)

Mach-Zender Interferometer Test

To obtain experimental values for the change in OPL over temperature, a Mach-Zender fiber interferometer was used. To measure the change in OPL vs. T, the change in temperature must be known, the length of fiber, the wavelength of light and the number of fringes or beats that occur over the temperature change. This set-up was extremely sensitive to vibration (as most interferometers are) and the temperature chamber had to be shut down and allowed to freely cool from a high temperature. Without a beat counter the most convenient approach turned out to be to measure the beat cycles per second and the change in temperature vs. time. The temperature slope of the fiber was assumed to be the same as the air in the chamber since the temperature change was slow.

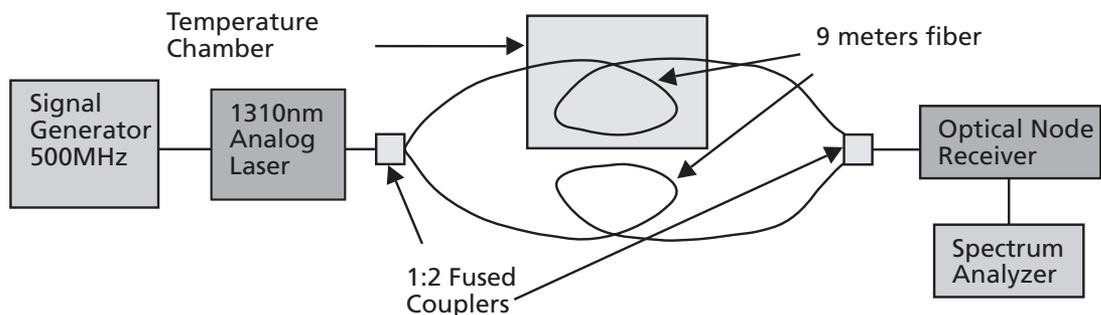


Figure 18: Experimental set-up for optical path length variation

Results

Seven measurements were made: 9.6, 8.3, 7.2, 10.2, 10.0, 7.2, 8.0 parts per million per degree centigrade. These average to 8.6 ppm/°C. This is closer to the ~7ppm/°C quoted by Passave Networks in a May 2001 IEEE 802.3ah presentation. The difference between these numbers and the Corning and Oriel numbers may be due to selection of the refractive index change at a wavelength other than 1310nm.

Analysis of Delay Change Effects on Upstream Synchronization

Estimate of Temperature Induced Delay per unit Time

We need to know the delay changes per unit time to understand synchronization implications. Light in optical fiber traverses 1 kilometer in 5 microseconds. Therefore the delay change is:

$$\Delta T/\text{°C} = 8.6 \times 10^{-6} \text{°C} * 5 \mu\text{s}/\text{km} = \mathbf{43\text{ps}/\text{°C}}$$

For the DOCSIS maximum length of 200miles (320km) of fiber, the temperature-induced jitter is:

$$\Delta T/\text{°C} = 43\text{ps}/\text{°C} * 320\text{km} = \mathbf{13.7\text{ns}/\text{°C}}$$

Finally, in order to understand the impact of this jitter on the need for periodic ranging in a fully synchronized system, we have to obtain a value for the maximum rate of temperature change for the fiber. From previous field data, a value of one degree centigrade per minute has been used. This should be conservative and includes temperature changes due to environment (solar load change from clouds to sun) and rapid cooling from rain on a solar heated cable. If we use this value we obtain:

$$\Delta T/T = 13.7\text{ns}/60\text{s} = \mathbf{0.23\text{ns}/\text{s}}$$

Delay Change due to Wind Loading

Aerial cable does stretch with wind. The construction of optical cable makes it tolerant of wind loading due to the loose tube construction that isolates the fiber from cable loading. Wind loading will affect aerial coaxial cables.

Due to the complexity and randomness of wind loading over long spans of coaxial cable, it is not clear if the loading effects will average out, reducing the possible peak values. It is difficult to analytically model this effect. Ideally, field measurements are a much more reliable method of investigating this. The following information on wind loading was an excerpt from the first published version of the DOCSIS 2.0 specification, version SP-RFlv2.0-I01-011231.



**Excerpt from Appendix VIII in the DOCSIS 2.0 RFI specification,
v. SP-RFlv2.0-I01-011231**

Wind loading is a difficult to deal with analytically because it is unlikely to be uniform along the cable. A delay model using a significant body of measured data is needed to investigate this further. Wind loading may be a source of fast delay variation and the ranging mechanism during station maintenance at the CMTS may not occur at intervals small enough to reduce this variation sufficiently.

The effects of wind loading on typical cable were investigated with a publicly available program from a coaxial cable manufacturer. These calculations showed that length changes in the range 0.01% and 0.05% are possible for various amounts of wind loading. This converts to significant propagation delay variation. As an example, with 5 miles (8 km) and 0.02% length variation, the change in propagation delay is:

$$(8/3e5)*(1/0.87)*2e-4 \text{ seconds} = 6 \text{ nanoseconds.}$$

This is a peak value, but the length of coax is quite short and the wind load is moderate. While the time duration over which this delay variation occurs is unspecified, it may be noted that wind gust data is readily available for most cities, and wind gust will be the primary mechanism for wind based timing changes on cable plants. For example, in New York City at the time of this writing, wind gusts of up to 40 mph are reported while average wind speed is about 10 mph. Hence, over a period of 1 to 4 seconds (the typical wind gust measurement interval), the wind speed changed by 30 mph. Much stronger wind gusts are frequently measured in locations prone to windy conditions.

Frequency of Station Maintenance

For Fiber-Temperature Induced Changes

If we wish to periodically range such that we adjust no more than 10% of the jitter requirement, then we must do periodic maintenance on each CM approximately once per second. If we are willing to allow the full movement of jitter allowance to the HFC plant, then periodic maintenance may be done once per 10 seconds. [A typical periodic ranging for an HFC system is once per 15 seconds. DOCSIS requires station maintenance about every 30 seconds (T4 time out has a maximum value of 35 seconds)].

For Coax-Wind Loading Induced Changes

The wind loading is potentially much more severe than the fiber temperature effects. There is not a time constant associated with the excerpt on wind loading above. However wind induced changes can be much faster than temperature induced changes. If the 6ns change mentioned above occurs on the order of a few seconds, then station maintenance may become a large portion of the upstream traffic.

Conclusions

It would seem that HFC delay variations due to temperature swings on very long aerial fiber lengths may be accommodated by performing station maintenance at a frequency around 1Hz. However there may be physical changes to the fiber length due to wind loading as well as temperature drift that may increase the frequency of station maintenance. It is most likely to show up in high QAM/Bandwidth channels. It will be important to obtain field measurements over a variety of HFC plants over a long time period to determine the impact of this problem.



REFERENCES

- [1] Error Control Systems for Digital Communication and Storage", Steven Wicker
- [2] "Comparison of Single-Carrier, Multi-Carrier, and Spread Spectrum Modulations for Upstream PHY Layer in HFC CATV Networks", by Jeyhan Karaoguz, John Yu, Vedat Eyuboglu, Motorola, IEEE 802.14a/98-018, July 13, 1998.
- [3] Prodan, R., et al., "Analysis of Two-Way Cable System Transient Impairments," CableLabs®, NCTA Proceedings, 1996.
- [4] Howard, D. and Roberts, H., "Dynamic Adaptation to Impaired RF Upstream Channels using Advanced PHY", Broadcom/ADC, NCTA Proceedings, 2002.

DOCSIS is a registered trademark of CableLabs®

Written by Hal Roberts
With contributions from Robert Jackson,
Benoit Legault, Michael Rude, Pawel Sowinski,
Dave Unger, Michael Wozniak and Jacey Godfrey



Web Site: www.adc.com

From North America, Call Toll Free: 1-800-366-3891 • Outside of North America: +1-952-938-8080 Fax: +1-952-946-3292
For a complete listing of ADC's global sales office locations, please refer to our web site.

ADC Telecommunications, Inc., P.O. Box 1101, Minneapolis, Minnesota USA 55440-1101
Specifications published here are current as of the date of publication of this document. Because we are continuously improving our products, ADC reserves the right to change specifications without prior notice. At any time, you may verify product specifications by contacting our headquarters office in Minneapolis. ADC Telecommunications, Inc. views its patent portfolio as an important corporate asset and vigorously enforces its patents. Products or features contained herein may be covered by one or more U.S. or foreign patents.

100606PR 04/02 Original © 2002 ADC Telecommunications, Inc. All Rights Reserved An Equal Opportunity Employer

