DYNAMIC ADAPTATION TO IMPAIRED RF UPSTREAM CHANNELS USING ADVANCED PHY

Daniel Howard Broadcom Corporation
Hal Roberts ADC

Abstract

The use of advanced receiver processing and system adaptation in the cable modem termination system (CMTS) can easily quadruple the upstream capacity by opening up new RF spectrum and by more efficiently using existing RF spectrum. The advent of DOCSIS 2.0 provides a 'toolkit' of physical layer features that allow this potential. However, since the upstream spectrum in previously unused portions of the band is highly dynamic in the level and type of interference present, it is critical that the CMTS be able to dynamically sense and adapt to changing channel conditions. Such dynamic adaptation ensures that the channel remains active, even in the presence of strong interference, and (as importantly) ensures that as the channel conditions improve, the capacity is restored to higher levels. The DOCSIS 2.0 'toolkit' provides many more options for the CMTS to handle ingress while maintaining high bandwidth, but without well designed adaptive algorithms the toolkit will be unused or worse, ill-used.

In this paper, an intelligent CMTS with advanced receiver processing and advanced system algorithms for dynamic adaptation is shown to provide significant benefits to existing deployments of DOCSIS 1.0 and 1.1 cable modems, as well as set the stage for future improvements using DOCSIS 2.0 technology. The performance improvements will be demonstrated in the presence of all of the most common upstream plant impairments: additive white Gaussian noise (AWGN), ingress, common path distortion (CPD), and impulse/burst noise. Mitigation of these impairments will be shown to open up spectrum below 20 MHz that may previously have been considered unusable. Further, the reliability of interactive services is increased by such dynamic adaptation, improving the market appeal of applications such as voice over IP over cable. Since the CMTS cost per subscriber is far less than either the cost of the cable modem itself and/or further plant upgrades, the solution described in this paper provides the lowest cost and fastest time to market approach for quadrupling the upstream capacity of existing cable modem networks.

INTRODUCTION

The spectral efficiency of current cable modem RF upstream channels can be greatly increased now that advanced PHY technologies such as higher orders of modulation, increased error correction, interleaving, ingress cancellation processing and better equalization are readily available in modern CMTS systems. Twice the spectral efficiency can be obtained by operating at 16 QAM instead of QPSK, more if the Reed Solomon (RS) forward error correction (FEC) overhead can be reduced. Three times the spectral efficiency can be obtained if new advanced PHY modems are deployed which can operate at 64 QAM. Further, RF spectrum that was previously unusable due to ingress, impulse/burst noise, or lack of sufficient equalization can now be used by all modems on the network, thereby creating even more capacity on the upstream for cable operators.

To obtain the greatest increase in spectral efficiency on the RF upstream, cable operators will have to deploy both advanced PHY cable modems as well as an advanced PHY Cable Modem Termination System (CMTS). However, if current plants are operating at QPSK levels on the upstream, the greatest incremental increase will come from increasing the signal constellation from QPSK to 16 QAM, and opening up new RF spectrum that was previously unusable. This incremental increase can be obtained merely by upgrading the CMTS since current DOCSIS 1.x modems can operate at 16 QAM as well as QPSK, but may have previously been unable to use 16 QAM either due to ingress or due to lack of sufficient equalization. Further, the most likely deployment scenario is a mixture of 1.x and advanced PHY modems since significant numbers of 1.x modems have already been and are still being deployed. Hence, the question becomes how to use an advanced PHY CMTS with a mixture of current and future advanced PHY modems to obtain the most capacity out of the cable network without requiring redeployment of modems or plant upgrades. This is the question addressed by this paper, and the answer turns out to be via adaptation of the mixed cable modem network to the various RF impairments that exist on the network.

The paper begins with a description of advanced PHY features, with emphasis on those that apply to DOCSIS 1.x modems as well as advanced PHY modems. Next, a discussion of channel impairments and mitigation strategies will be used to show how an adaptive CMTS using advanced PHY features can keep channel capacities high most of the day, and only increase robustness (thereby dropping capacity) when the channel conditions require it. To accomplish this, an adaptation system is described which incorporates spectrum monitoring with control of CMTS burst profiles for the various traffic types being transported. The result is operation at high spectral efficiencies for the great majority of the day, meaning more throughput to users or alternately more users possible on current upstream channels.

ADVANCED PHY FEATURES AND APPLICATION TO DOCSIS 1.X MODEMS

As described in a companion paper in this session [1], DOCSIS 2.0 advanced PHY technology includes both advanced TDMA (ATDMA) and synchronous CDMA (SCDMA). These advanced PHY technologies increase the capacity in clean upstream channels by providing higher spectral efficiencies, with up to 64 QAM for ATDMA in the specification, (and up to 256 QAM in some vendor implementations). Advanced PHY also provides significant increases in the robustness of upstream signaling against the most common RF impairments: additive white Gaussian noise (AWGN), ingress of radio/navigation signals, common path distortion (CPD), and impulse/burst noise. For a detailed analysis and modeling of upstream impairments, the reader is referred to a previous NCTA paper by one of the authors, which also has references to other upstream measurements and modeling papers [2].

However, there are several features in advanced PHY CMTS systems that improve not only performance with advanced PHY CMs, but also the performance of existing 1.x CMs. A listing of several advanced PHY features, some of which apply to 1.x CMs is shown in Table 1.

Table 1. Advanced PHY Features

Feature	Improves 1.x ?
Improved AWGN Mitigation	
Lower implementation loss	YES
Better receive equalization	YES
Improved burst acquisition	YES
More FEC (T=16)	NO
Ingress/CPD Cancellation	
Cancellation of ingress	YES
Cancellation of CPD	YES
Improved Mitigation of Impulse and Burst Noise	
Cancellation of spectral peaks in periodic impulse noise	YES
More FEC (T=16)	NO
RS byte interleaving	NO
SCDMA mode	NO

As is seen from the table, existing DOCSIS 1.x modems will benefit significantly from deployment of a modern advanced PHY CMTS in mitigation of AWGN, ingress/CPD, and impulse/burst noise. The most significant improvement for 1.x modems is in the area of cancellation of ingress and/or CPD: over 30dB of narrowband ingress can be cancelled from the received spectrum. Since all of the processing for cancellation resides in the new CMTS, existing modems of any DOCSIS version will benefit from the advanced PHY technology. Figure 1 shows the improvement possible in 16 QAM mode against multiple ingressors, and Figure 2 shows the improvement possible against wideband ingressors. The latter impairment case is also applicable to a group of ingressors that must be cancelled as a zone as opposed to individual cancellation.

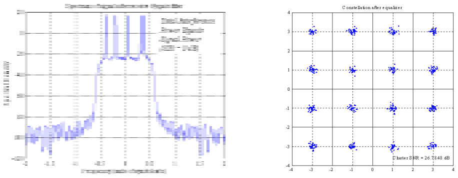

Figure 1. Cancellation of 5 Ingressors in 16 QAM mode, SIR=0 dB

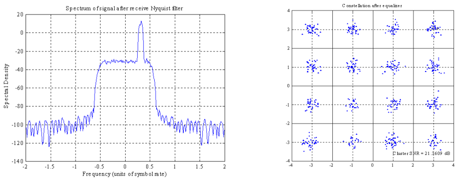

Figure 2. Cancellation of Wideband Ingress in 16 QAM mode, SIR=0 dB

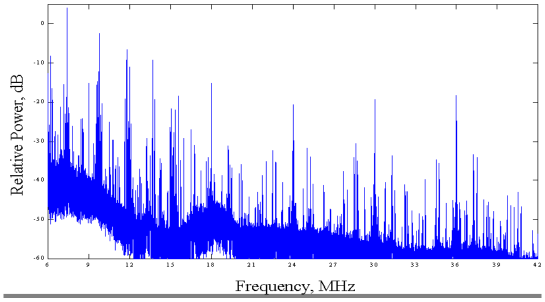

Figure 3. CPD may be seen as multiple narrow band ingressors, capable of being cancelled by an adaptive ingress filter.

The impact of the ability to cancel ingress and CPD from 1.x modems is significant: DOCSIS 1.x CMs can now be operated with less FEC in channels heavy with ingress and/or CPD, such as channels below 20 MHz which may have previously been considered unusable.

Further, the improved equalization and the lower implementation loss in advanced PHY CMTS hardware means that ingress-free channels, which previously could not support 16 QAM, can now easily support it.

The improvement in mitigation of impulse/ burst noise is less substantial than that of ingress and CPD, but nonetheless exists. For example, a previous paper by one of the authors showed that periodic impulse noise could be tracked and avoided by intelligent scheduling in an advanced CMTS. Note also that most impulse noise has a non-uniform spectral signature. This characteristic may be exploited and significant amounts of the noise energy reduced via the ingress cancellation processing. If either of the latter capabilities are not supported in the advanced PHY CMTS, the benefits of advanced PHY can still be reaped via deployment of data-only service in channels which have high impulse/burst noise. The additional packet loss due to impulse noise can easily be low enough for the most commonly occurring events that the result is a degradation that most users would seldom notice.

Thus, an advanced PHY CMTS can result in current DOCSIS modems operating at 16 QAM, and additional spectrum below 20 MHz being usable by these modems. If data-only modems are moved below 20 MHz, there would be more room at higher frequencies for services requiring higher quality of service or higher channel capacities.

EVOLUTION OF DEPLOYMENT STRATEGIES

In the previous section, the notion of how MSO's could deploy advanced technology was introduced, especially as it relates to the existence of legacy systems on the plant. In this section, a more detailed look at deployment strategies is presented; in particular, the transition from *hardware-limited* capacity to *spectrum-limited* capacity is described. Given the constraint of legacy modems on the plant and the desire to use spectrum below 20 MHz, it will be seen that the logical transition path is from upstream channels with relatively fixed center frequencies, symbol rates, and modulations to channels which can adapt the modulation, center frequency, and symbol rate to the instantaneous conditions of the channel.

Currently, many MSOs have more RF bandwidth available than upstream receivers, as evidenced by node combining and the lack of utilization of all RF upstream spectrum on the plant. Since often the modulation scheme currently in use is QPSK with maximum levels of FEC, the only adaptation strategy available is to hop the cable modem frequency. With the advent of guaranteed pre-equalization in DOCSIS 1.1, 16-QAM may be used for the majority of spectrum. Note that the fact that that the upstream can support high QAM levels is evidenced by successful deployments of ADCs OFDMA based cable telephony system in 32-QAM mode. OFDMA is inherently robust against multipath, however multipath may be mitigated by pre-equalization in TDMA/SCDMA systems.

INCREASING UPSTREAM BANDWIDTH

Already some MSOs have run out of upstream bandwidth due to a variety of factors: low frequency ingress, legacy FDM set top boxes, and bandwidth allocated to other entities, including government agencies and schools.

At the same time, due to the advent of more symmetric services, such as 'Voice Over IP', the demand for upstream bandwidth is beginning to increase. The MSOs are either currently faced with or soon will be facing the choice of improving the usage of the limited upstream resource, or to engage in expensive plant upgrades or node splitting.

As we shall see, the improved modulation techniques of advanced PHY systems opens the potential of optimally using this upstream resource.

Before utilizing advanced PHY, for many MSOs the first step to increasing upstream capacity is node 'decombining'. In systems with lower penetration rates, a common practice has been to combine the upstream signals from multiple node into the same CMTS receiver. As penetration rises, decombining the upstream nodes results in a low cost upstream bandwidth expansion. This allows separate optical nodes to be serviced by separate CMTS channels. To achieve this it is necessary to evolve to higher density CMTS receivers. Ultimately, however, if costs are to be contained it is necessary that the CMTS utilize the spectrum assigned to it in the most efficient manner.

MAXIMUM SPECTRAL EFFICIENCY IN DYNAMIC IMPAIRMENTS

As was stated above, the DOCSIS 1.x and 2.0 PHY 'toolkits' provide the potential for maximum utilization of the upstream channel. However, this cannot be realized without intelligent CMTS sensing, analysis and reaction mechanisms.

In addition, it will be necessary to have spare upstream receivers in the CMTS to take advantage of 'divide and conquer' strategies that will be seen below.

A well-known technique used in DSL (digital subscriber loop) to achieve "Shannon limited" capacity in available channels is to slice up the available bandwidth into micro-channels and adjust the modulation parameters to the maximum bandwidth that the micro-channel conditions will allow. With the use of spare upstream CMTS receivers, this same technique may be used in a coarser fashion to maximize the capacity of HFC upstream bandwidth (see Figure 5 below).

Dynamic Adaptation

Obviously dynamic adaptation only makes sense if the channel is changing dynamically. That this is true of HFC upstream channels is evidenced by many studies [3]. These studies show that impulse and burst noise is often higher at lunch and dinner times. Ingress is often higher at night, while CPD varies with temperature, humidity and wind due to the source of CPD, cable connectors. (CPD, or Common Path Distortion, is often caused by a diode effect resulting from the corrosion of cable connections.)

These same studies show that the percentage of time that impairments exist on the plant is usually well below 10%. While the new modulation technologies embodied in DOCSIS 2.0 provide increased robustness to tolerate these higher levels of impairments, the increased robustness may be had at the expense of spectral efficiency. On large packets for example, the spectral efficiency can be reduced from the maximum of about 4.7 bits/Hz (64 QAM, minimum FEC) to about 1.2 bits/Hz (QPSK, maximum FEC), a 4x reduction in capacity (see Figure 4 below). On a plant with 12 MHz of bandwidth currently used for upstream data service, the capacity can thus be varied from 14 Mbps to 57 Mbps with DOCSIS 2.0 modems on the plant. Clearly, dynamic adaptation is a key strategy in the optimization of upstream bandwidth.

The variation is slightly less when the modems on the plant are all DOCSIS 1.x, but is still dramatic. The maximum capacity would result from using 16 QAM, which has a spectral efficiency of about 3 bits/Hz. In this case, the plant capacity can be varied from 14 Mbps to about 36 Mbps. Further, dynamic adaptation can also open up new spectrum for service and increase capacity. If 6 MHz below 20 MHz were useable at 16 QAM with advanced CMTS processing most of the day, this would give an additional 18 Mbps of upstream capacity with existing DOCSIS modems. Thus, by upgrading the CMTS to one that uses dynamic adaptation to leverage ingress cancellation, improved equalization, burst acquisition and lower implementation loss, existing DOCSIS 1.x networks can be expanded from 14 Mbps to 54 Mbps, again a four-fold increase in capacity.

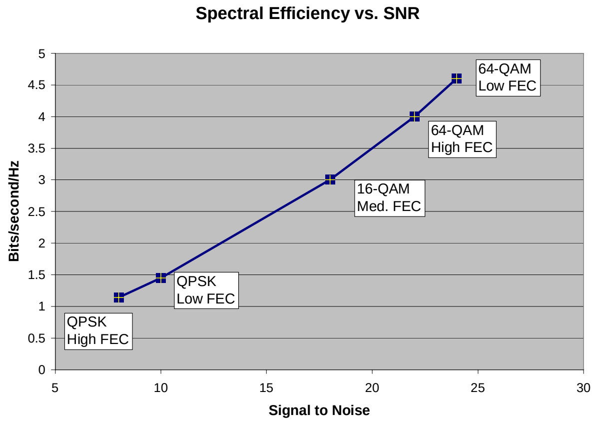

Figure 4. Spectral Efficiency vs. SNR

ADAPTATION TECHNIQUES

Adaptive Modulation

The examples of capacity improvements in the previous section point to simple adaptation techniques. First, the level of FEC used on packet transmissions can be increased as impairments increase. Over 6 dB improvement in robustness is possible with this technique, albeit with a 20-30% drop in spectral efficiency. The other simple technique is to decrease the order of modulation. As previously discussed, with 2.0 modems, 64QAM can be reduced to 16 QAM for 6 dB of improvement in robustness, or all the way down to QPSK for 12 dB of improvement. Taken together, changing FEC and modulation order provide up to 18 dB of improvement in robustness at the expense of 75% of network capacity. As can be seen by the plots of ingress from 24 nodes below, the ingress noise appears to vary from one end of the spectrum to the other by 15 dB to 20 dB, confirming that spectrum may be opened up by using the modulation choices available in advanced PHY.

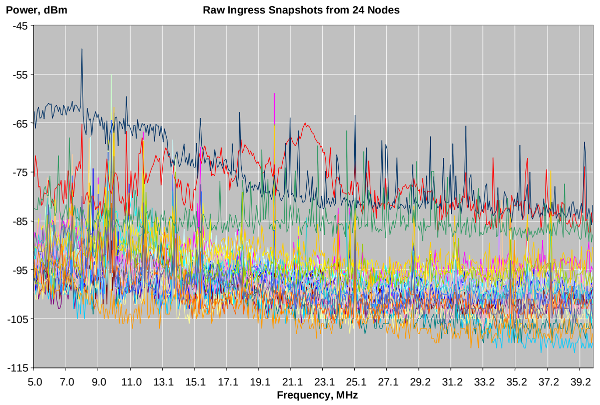

Figure 5. Upstream Plots from 40 Optical Nodes shows Ingress Decreasing by 15 dB to 20 dB at Higher Frequencies.

Channel Hopping

The next level of adaptation involves changing the channel center frequency to avoid significant levels of impairments such as ingress. While the availability of ingress cancellation technologies in the CMTS will reduce the necessity of changing channels much of the time, this adaptation technique remains viable for MSOs with spare RF upstream bandwidth. However, since now only the highest levels of ingress need be avoided, the frequency hop adaptation technique may now involve slight shifts in carrier frequency as opposed to hopping to an entire new block of spectrum on the upstream.

Decreasing Symbol Rate

Next, the symbol rate can be reduced for increased robustness against all types of impairments, again at the cost of reduced capacity. Assuming the transmit power is maintained at the original level, a reduction in symbol rate by a factor of 2 will add 3 dB more robustness against AWGN, ingress, and burst noise. Further, the length of an impulse event that can be corrected is doubled by the fact that in the time domain, the symbols are twice as long as before, therefore fewer symbols are corrupted by the same impulse event. Both aspects of symbol rate reduction are shown in Figure 6.

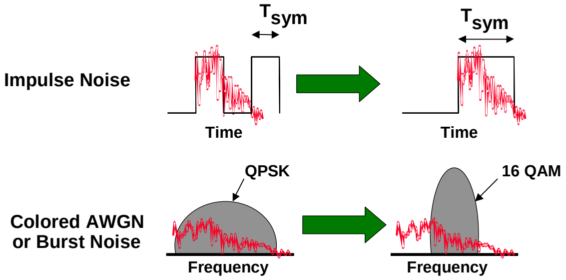

Figure 6. Adaptation via Symbol Rate Reduction.

Thus, for MSOs with spare RF spectrum but without spare upstream receivers, the following adaptation techniques apply:

Increase FEC
Reduce order of modulation
Reduce symbol rate
Change carrier frequency

If, on the other hand, the operator has run out of available bandwidth, and wishes to avoid costly plant upgrades, an effective next step is to employ backup CMTS receivers in the adaptation process. In particular, if the symbol rate is reduced to mitigate impulse/burst noise, the capacity on the plant can be maintained by dividing the channel into smaller subchannels with the same spectral power density. Thus the symbol rate and center frequency must be changed for this adaptation technique. Note that when altering the center frequency of current upstream channels, the following conditions apply:

1) Reranging is generally required.
2) If the order of modulation is already reduced to QPSK, reranging will likely not be required as pre-equalization can be avoided.

The benefits of channel dividing against narrowband and broadband impairments are shown in Figure 7 below, where it is seen that dividing an RF channel which previously could only support QPSK produces subchannels which support much higher orders of modulation. The technique merely requires the availability of backup upstream receivers to optimize the capacity of the network under impaired conditions.

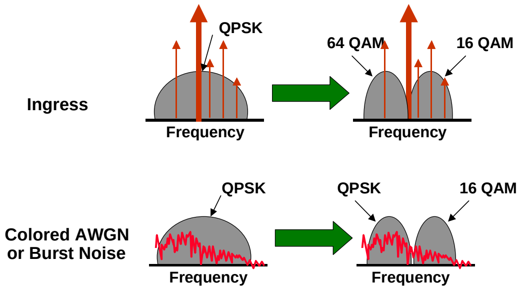

Figure 7. 'Divide and Conquer Strategy' - Channel Dividing to Combat Ingress and Broadband Noise.

Note that special considerations exist for mixed DOCSIS 1.x/2.0 channels. The CMTS must not select an adaptation technique that only works for 2.0 modems, although 2.0 specific techniques can be applied to the 2.0 modems as long as an alternative for the 1.x modems is applied as well. For example, if moderate impulse noise is detected, the CMTS could increase the interleaving on 2.0 modems while maintaining the order of modulation at 64 QAM, and reduce the order of modulation on 1.x modems. Alternatively, the 2.0 modems could switch to SCDMA mode, if impulse conditions warrant. If the impulse noise is too long for simple constellation changes, the symbol rate of all modems on the network may need to be reduced so that the 1.x modems stay active. There are also differences in the equalization capabilities of 1.x and 2.0 modems, and this may also lead to a different adaptation strategy when mixed networks are deployed.

Finally, additional adaptation techniques will likely exist depending on vendor implementations. For example, a smart scheduling alternative to periodic burst noise exists if the noise can be detected and tracked. In this case, the packets from 1.x modems (and the 2.0 modems if necessary) can be scheduled around the impairment without requiring a symbol rate reduction. And the modems may similarly have vendor specific performance and/or adaptation capabilities. Hence, the rules for adaptation should take into account any and all differences in CMTS and modem capabilities.

SYSTEM ADAPTATION

The heart of any scheme to dynamically adapt to changing channel impairments is the ability to detect and classify RF impairments on the upstream. Figure 8 depicts a basic adaptation system, with key components being the spectrum monitor and a lookup table of burst profiles. The spectrum monitor can be internal or external to CMTS, but it is important that RF impairment detection and classification process use rules based on plant measurements and impairment models, such as those presented in [2].

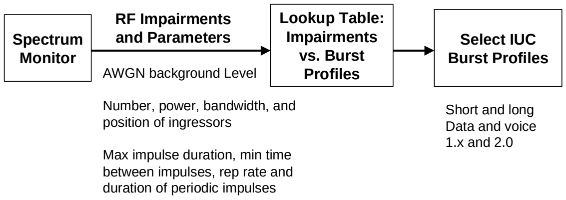

Figure 8. Adaptation Process.

In particular, the spectrum monitor should classify each impairment separately, since different adaptation strategies exist for different impairments. For example, if the total interference power used to characterize channel, then ingress cancellation and FEC/interleaving will not be leveraged to their fullest extent. Consider the case with an AWGN background noise floor that is 22 dB down from the signal power level, but an ingress signal is present that is 10 dB above the signal power. A 2.0 modem could easily operate at 64 QAM and a 1.x modem could operate at 16 QAM in this level of noise. But if the total interference power were used to characterize the channel, the system would erroneously assume the channel was unusable due to SNR being too low for even QPSK operation.

Further, the spectrum monitor should examine both in-band and out-of-band impairments to be most effective. In the case of a single strong in-band ingress signal that is near the channel edge, a slight shift of center frequency only may be required to keep the channel active and at peak capacity. If the symbol rate is to be reduced without the creation of additional subchannels, the best position for the signal with the reduced symbol rate must be determined. Finally, for impulse/burst noise adaptation, the spectrum monitor should also have the capability to measure impairments in the time domain as well as in the frequency domain.

Once the RF impairments have been detected and classified, the results must be used to determine the burst profiles for the channel that optimize capacity while maintaining sufficient robustness against the impairments. A lookup table is one approach to this requirement, where the system performance is characterized in a lab against a variety of impairments and levels, and optimum burst profiles determined for each impairment and level of impairment. Table 2 shows an example lookup table for AWGN impairments with coarse parameter changes. In reality, a finer table would be desired so that the burst profiles can truly be optimized for the channel conditions that exist.

Table 2. Lookup Table for AWGN

SNR	Modulation	FEC
35 dB	256 QAM	Low[1]
30 dB	256 QAM	Med
25 dB	64 QAM	Low
20 dB	64 QAM	High
15 dB	16 QAM	High
10 dB	QPSK	Med

As the FEC overhead is increased and the modulation type reduced, the spectral efficiency will drop, but for the benefit of greater robustness. The actual FEC used in the burst profile will depend on the packet size, quality of service required, and so on. For example, one set of tables could apply to a packet error rate of less than 1%, while another set of tables could allow error rates of up to 5%. The former could then be applied to voice packets and the latter to best effort data packets. Hence, there could be several lookup tables for each type of service and packet size that optimizes the burst profile subject to the main constraint of tolerating the given level of AWGN with a selected packet error rate.

Similar lookup tables can be developed for each impairment and even combinations of impairments. In this manner, when any previously seen (or postulated) combination of impairments are detected on the cable upstream, the CMTS can use the optimum burst profiles for those particular impairments.

Customization of Algorithms

The algorithms described above, which are embedded in the CMTSs, should allow customization by the MSO. This is due to the fact that HFC plants may differ greatly in the typical ingress signature. Plants with large impulse noise compared to ingress may require different optimization than the reverse situation. In addition, algorithms will need modification over time as experience grows and/or the mix of 1.x/2.0 modems migrates towards all 2.0 modems.

Intelligent CMTS Initialization

How might such a system operate in a global sense? Upon boot-up, the CMTS would characterize the upstream channel using the spectrum monitor. Next, an initial burst profile based on the detected AWGN background level could be selected. Again, ingress and impulse/ burst noise power in the channel must not be used for this decision if it is to be optimal.

Next, if ingress is present and levels are too high for cancellation in the measured AWGN background, a new burst profile can be selected that the ingress canceller can handle. The same process can be used for other impairments such as impulse/burst noise, although some adaptation techniques such as interleaving are fairly independent of the background AWGN and ingress power levels.

CONCLUSIONS AND FUTURE DIRECTIONS

The need for, benefits of, and basic design aspects of an adaptation system for mixed DOCSIS cable modem networks have been described in this paper. It was shown that adaptation can quadruple the bandwidth on the network during the great majority of the day. Adaptation also reduces the capacity during detected impairments but in a manner that keeps spectral efficiency as high as possible for the detected impairments and returns to the highest spectral efficiency when the impairment diminishes.

As MSOs transition from hardware limited capacity to spectrum limited capacity, more complex adaptation schemes will be employed, for example channel dividing. This is

---

[1] Proprietary 256-QAM Mode

necessary to adapt 1.x modems while maintaining capacity, but as more 2.0 modems are deployed, the need for channel dividing may be reduced. Further improvements in the robustness and capacity of DOCSIS modems can also lead to modifications of the strategies described here.

REFERENCES

[1] Legault, B., “A Technical Analysis of DOCSIS 2.0,” NCTA Proceedings, 2002.

[2] Howard, D. “Detection and classification of RF impairments for higher capacity upstreams using advanced TDMA,” NCTA Proceedings, 2001.

[3] Prodan, R., et al., "Analysis of Two-Way Cable System Transient Impairments," CableLabs®, NCTA Proceedings, 1996 and "Two-Way Cable Television System Characterization", Cable Television Labs, 1995.