AN ANALYSIS OF THE TDMA AND S-CDMA TECHNOLOGIES OF DOCSIS 2.0

Fabien Buda, Emmanuel Lemois, and Hikmet Sari
Juniper Networks
4-14 rue Ferrus, 75014 Paris, France

Abstract

This paper analyzes the performance of TDMA and S-CDMA, which are the two modula-tion and multiple access techniques included in the newly developed DOCSIS 2.0 specifications. The two techniques are analyzed in the presence of linear signal distortion, ingress noise, and burst noise, which characterize cable plants. We also discuss how the transmission parameters can be adopted in both multiple access techniques to plant conditions in order to increase capacity or robustness. The results give a good summary of the relative merits of TDMA and S-CDMA.

1. INTRODUCTION

Data over Cable System Interface Specifications (DOCSIS) [1] elaborated under the leadership of Cable Television Laboratories, Inc. (CableLabs) have established themselves as the major industry standard for two-way communications over hybrid fiber coax (HFC) cable plants. The first specifi-cations, referred to as DOCSIS 1.0, have been largely deployed in cable networks, and the second specifications (DOCSIS 1.1) are now in the early phase of deployment. The physical (PHY) layer of DOCSIS 1.1 specifications is the same as that of DOCSIS 1.0, the difference between the two sets of specifications lies on the medium access control (MAC) layer, which includes a quality of service (QoS) in DOCSIS 1.1 specifications.

The demand for increased capacity on the upstream channel has stimulated CableLabs to draft a new set of RF interface specifications referred to as DOCSIS 2.0. These specifications, which were developed with the vendor community, also aim at increasing the robustness to various impairments on the upstream channel in cable plants. The basic features of the new specifications include an increased channel bandwidth (6.4 MHz), several additional modulation schemes including 64-QAM which gives a 50% increase of the throughput with respect to the 16-QAM modulation that is in the current DOCSIS 1.0 and 1.1 specifications, and an improved forward error correction scheme.

One of the most controversial issues during the elaboration of the DOCSIS 2.0 specifications was the multiple access scheme, which describes how different cable modems (CMs) access a particular physical channel. The two techniques in compe-tition were the conventional time-division multiple access (TDMA) used in DOCSIS 1.0 and 1.1 as well as in many other international standards, and the synchronous code-division multiple access (S-CDMA) scheme used in some proprietary systems. The decision was eventually made to include both techniques in the specifications, and make both of them mandatory for both the CM side and the cable modem termination system (CMTS) side.

The two multiple access technologies included in DOCSIS 2.0 have very basic differences in terms of their operating principles and robustness to channel impairments and equipment imperfections. This paper gives the results of an exhaustive perfor-mance evaluation of DOCSIS 2.0 TDMA and S-CDMA in hybrid fiber/coax (HFC) cable networks. The evaluation is performed in the presence of most common impairments in cable plants which include linear signal distortion, additive white Gaussian noise (AWGN), ingress noise (characterized as narrowband interference) and burst noise. We also discuss how the transmission parameters can be adopted to plant conditions in order to increase capacity or robustness.

The paper is organized as follows: In Section 2, we give a brief review of DOCSIS 2.0. Section 3 discusses the noise performance, and more particu-larly the trade-off between data rate and perfor-mance, and the optional trellis code in the S-CDMA mode. In Section 4, we study the influence of channel impairments including linear distortion, ingress noise, and burst noise. Finally, we give our conclusions in Section 5.

2. A BRIEF DESCRIPTION OF DOCSIS 2.0

Working with the vendor community, CableLabs has recently released the interim DOCSIS 2.0 RF interface specifications [2], which aim at increasing the capacity and robustness to various impairments of the upstream channel in cable plants. These specifications do not affect the downstream channel, neither do they affect the medium access control (MAC) functions except for the changes required to accommodate the new physical (PHY) layer. The basic features of the new specifications include an increased channel bandwidth (6.4 MHz), several additional modulation schemes including 64-QAM which gives a 50% increase of the throughput with respect to the 16-QAM modulation that is in the current DOCSIS 1.0 and 1.1 specifications, and an improved forward error correction scheme.

The previous DOCSIS specifications (DOCSIS 1.0 and 1.1) were based on time-division multiple access (TDMA). More specifically, time-division multiplexing (TDM) is used on the downstream channel (from CMTS to CMs) and TDMA is used on the upstream channel (from CMs to CMTS). For DOCSIS 2.0, two multiple access proposals were made. One of these is to keep the TDMA used in previous DOCSIS specifications, and the other is synchronous code-division multiple access (S-CDMA), previously used in some proprietary modems. The decision was made in August 2001 to include both schemes in DOCSIS 2.0.

2.1. TDMA

TDMA is a simple and popular multiple access technique used today in many international stan-dards and proprietary systems. It consists of assig-ning different time slots to different users. During the time slot assigned to one user, all other CMs remain silent, and therefore, there is no interference between users. Since DOCSIS 1.x is based on TDMA, the use of this technique in DOCSIS 2.0 is natural and its implementation requires little effort. Furthermore, since DOCSIS 2.0 requires backwards compatibility with DOCSIS 1.x, implementation of TDMA in DOCSIS 2.0 equipment is unavoidable.

The upstream spectrum in cable networks extends from 5 MHz to 42 MHz (to 65 MHz in Europe). The upstream channel width in DOCSIS 1.x is 2^{n-1} times 200 kHz, with n = 1, 2, 3, 4, 5. DOCSIS 2.0 adds a 6.4 MHz channel bandwidth corresponding to n = 6. In all cases, the raised-cosine Nyquist roll-off factor is 0.25, and the symbol rate is 0.8 times the channel bandwidth. A number of upstream channels may be available in the 5 – 42 MHz upstream spectrum depending on what part of this spectrum is used for legacy (analog or non-DOCSIS digital) channels and what part of it is not usable due to excessive noise, interference, or distortion. TDMA on the cable upstream channel is actually a combination of frequency-division multiple access (FDMA) and TDMA. When a CM makes a resource request, the CMTS grants time slots on one of the upstream channels, each of which accommodates a number of CMs in the TDMA mode. That is, the multiple access scheme is actually

FDMA/TDMA, but it is referred to as TDMA for simplicity.

The upstream modulation schemes in DOCSIS 1.x are the quaternary phase-shift keying (QPSK) and the quadrature amplitude modulation (QAM) with 16 constellation points (16-QAM), which doubles the data rate. To these, DOCSIS 2.0 adds 8-QAM, 32-QAM, and 64-QAM, but only the latter modu-lation is mandatory at the CMTS. With respect to 16-QAM, this modulation increases the data rate by 50% but loses 6 dB in signal-to-noise ratio (SNR), defined as the transmitted energy per symbol to the noise spectral density ratio (E_s/N_0). The inclusion of 8-QAM and 32-QAM in the specifications allows reducing the granularity in the trade-off between data rate and performance.

Since DOCSIS 2.0 also aims at increasing robust-ness, it extends the error correction capacity of the Reed-Solomon (RS) code from 10 to 16 bytes per block. Furthermore, a byte interleaver is included in the TDMA specifications so as to break the error events caused by burst noise and uniformly distri-bute the resulting symbol errors. The interleaver is a block interleaver whose length is equal to the RS block length and depth is a configurable parameter, which distributes the error bursts over a selected number of RS blocks.

2.2. S-CDMA

The S-CDMA of DOCSIS 2.0 is actually not a true CDMA, but rather a mix of code-division multip-lexing (CDM), CDMA, and TDMA. The incoming data is organized in mini-slots, which have two dimensions (spreading codes and time). The time duration of mini-slots is one S-CDMA frame that spans a programmable number of S-CDMA symbol intervals. (The maximum frame length is 32 S-CDMA symbol intervals.) Symbol spreading is performed through multiplication by a spreading code (spreading sequence) of 128 chips taken from a set of 128 orthogonal codes that are generated by a quasi-cyclic shift. Spreading is in the time domain, which means that an S-CDMA symbol interval is equal to 128 TDMA symbol intervals.

A mini-slot contains a programmable number of spreading codes, which can be as low as 2 and as high as 128. A mini-slot contains symbols from a single CM. Suppose that the number of codes per mini-slot is 4 and that the frame length is 16. Then, the mini-slot contains 64 symbols, and a given code is assigned to the same user for a time duration of 16 consecutive S-CDMA symbols. The 4 symbols transmitted in parallel within the same mini-slot in this example are code-division multiplexed. If all mini-slots of an S-CDMA frame are assigned to the same cable modem, S-CDMA is reduced to pure CDM during that interval.

To the other extreme, if the mini-slots contain two codes only, and each mini-slot is assigned to a different CM, then there is code-division multiple access between 64 CM signals during that frame. The spreading code orthogonality ensures that there is no interference between symbols transmitted in parallel by the same CM, since these symbols are perfectly time synchronized. But interference arises between signals generated by different CMs, due to non-ideal timing synchronization. To limit the resulting degradation, DOCSIS 2.0 specifies that the maximum timing error between a CM and the CMTS must not exceed 1% of the chip interval.

In addition to the RS code, S-CDMA specifications also include trellis-coded modulation (TCM) as an option. Trellis coding reduces the number of information bits per symbol by one, and so trellis-coded 64-QAM (referred to as 64-TCM) is equivalent to uncoded 32-QAM in terms of spectral efficiency. Therefore, S-CDMA specifications also include 128-TCM, which is strictly equivalent to uncoded 64-QAM as far as information bit rate is concerned. But the S-CDMA specifications neglect to include an interleaver between the external RS code and the internal trellis code, which reduces the benefit of trellis coding. In contrast, S-CDMA specifications include some interleaving after the trellis encoder to reduce the effect of burst noise. This interleaver operates on subframes, and interleaving is different for uncoded bits and coded bits. Subframes are independent of mini-slots, and a subframe is always contained within a single frame.

Symbol spreading in S-CDMA is used for the traffic mode only. That is, the spreader is turned off (S-CDMA is deactivated) during initial ranging and periodic station maintenance. The implication of this is that whatever this multiple access technique may offer with respect to TDMA, the benefit is restricted to the traffic mode.

3. NOISE PERFORMANCE

3.1. Flexibility of the Standard

DOCSIS 2.0 is a toolbox that gives full flexibility in trading off data rate against performance and robustness to channel impairments. Indeed, it includes all QAM signal constellations from 4-QAM (QPSK) up to 64-QAM, and also the RS code correction capability can take all values from $RSt = 0$ (no coding) up to $RSt = 16$, for different block lengths. First, it is well known that with respect to 2^m-QAM, the 2^{m+1}-QAM signal constellation increases the number of bits per symbol by one, but requires 3 dB higher signal-to-noise ratio (SNR) to achieve the same bit error rate (BER) performance.

This means that the constellation alone offers a trade-off between data rate and performance with a granularity of 3 dB in SNR or transmitted signal power. The only exception to this general rule is the step between 4-QAM (QPSK) and 8-QAM, which is approximately 4 dB. For example, if the cable network is operating with 64-QAM and the noise level increases, then switching to 32-QAM increases noise margin by 3 dB, and switching to 16-QAM increases it by 6 dB. In practice, it is desirable to have a finer granularity in this trade off. The second set of parameters that make this possible are the block length and error correction capacity of the RS code. Since there is a large degree of freedom in selecting the code parameters, the granularity of the trade off between data rate and performance can be made as fine as desired.

At this point, it is instructive to discuss the additional flexibility in the S-CDMA mode. As presented in Subsection 2.2, the total number of spreading codes is 128, but this number can be arbitrarily reduced in order to improve performance. This property of S-CDMA is actually often put forward as an advantage with respect to TDMA. While it is correct that S-CDMA allows this type of trade off, it is of little interest as it sacrifices too much data rate compared to what is offered by the modulation and code parameters. Indeed, reducing the number of codes from 128 to 64 in S-CDMA gives an SNR gain of 3 dB, but this is achieved by sacrificing 50% of the data rate. If we assume that the network is using 64-QAM, the same gain is achieved by switching down to 32-QAM, and this gain only involves a 16.6% of the data rate. In fact, the noise margin can be increased by 9 dB through modulation (by switching from 64-QAM to 8-QAM) if one is prepared to sacrifice 50% of the data rate.

Fig. 1 shows the spectral efficiency (number of information bits per symbol) vs. the SNR required for BER = 10^{-8} for different signal constellations, RS code parameters, and also the number of spreading codes in S-CDMA. Rather than giving some discrete points corresponding to a few values of the *RSt* parameter, this figure gives for each constellation a continuous curve corresponding to all *RSt* values from $RSt = 0$ to $RSt = 16$ with an RS block length of 255, and then to lower RS block sizes with *RSt = 16*. For uncoded 64-QAM and RS-coded QPSK, the figure also gives a curve that illustrates performance vs. spectral efficiency in S-CDMA with a reduced number of codes.

This figure clearly shows that for each signal constellation, there is a range of the number of information bits per symbol over which that constellation is the best to use. For example, the 64-QAM constellation is optimum down to 4.8 bits per symbol, and 32-QAM is optimum between 4.8 and 3.8 bits per symbol. It is also clear from this figure that reducing the number of spreading codes in S-CDMA to improve performance is not attractive as long as such a compromise can be made by the modulation and coding functions.

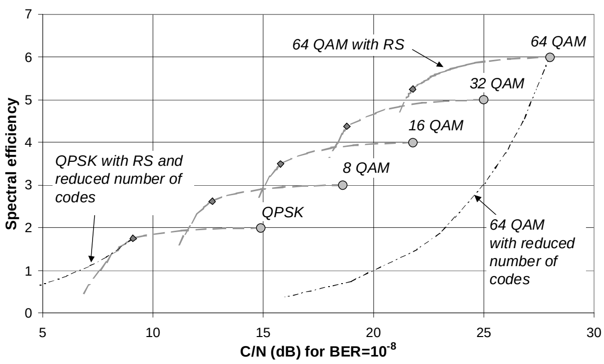

Fig. 1: SNR required with different signal constellations, RS code rates, and number of spreading codes in S-CDMA.

The only case where reducing the number of spreading codes in S-CDMA may be of interest is when the lowest level modulation (QPSK) together with *RSt = 16* in the RS code are not sufficient to get the desired performance. But even if this were a situation of strong practical interest, the current S-CDMA specifications would not help, because symbol spreading is disabled during initial ranging and periodic station maintenance, which are most critical to overall system performance.

3.2. Optional Trellis Coding

As mentioned earlier, the DOCSIS 2.0 specifica-tions include an optional TCM in the S-CDMA mode. The trellis code is an 8-state code, which gives an asymptotic coding gain of 4 dB. But what is more significant is the gain provided by the trellis code when cascaded with the mandatory RS code. Depending on the *RSt* parameter of the RS code, a BER of 10^{-2} - 10^{-4} at the TCM decoder output is reduced to less than 10^{-8} by the RS decoder. The asymptotic gain of the concatenated coding scheme is therefore the TCM coding gain somewhere in the range of 10^{-2} - 10^{-4}. With the trellis code used, this gain is 1.1 dB for *RSt = 16* and 2.5 dB for *RSt = 2*.

But the coding gain indicated above assumes an infinite interleaver between the RS encoder and the trellis encoder, so that the deinterleaver on the receive side can break the error events at the TCM decoder output and distribute them over a large number of RS blocks. The number of RS symbol errors per block is then small, and these errors can be located and corrected by the RS decoder. Unfortunately, the DOCSIS 2.0 specifications do not include such an interleaver but instead a sub-frame interleaver that keeps a part of the information bits uninterleaved. This significantly reduces the TCM gain, as illustrated in Fig.2 for 128-TCM. The RS code used in this figure is of length 255 and its correction capacity is *RSt = 16*.

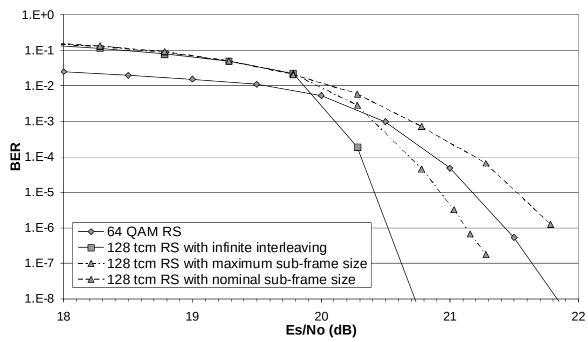

Fig. 2: BER performance of 128-TCM.

We can see that at the BER of 10^{-8}, 128-TCM gains 1.1 dB with respect to 64-QAM when it employs an infinite interleaver prior to the trellis encoder. With the subframe interleaver of DOCSIS 2.0, the gain is a function of the subframe size. With the maximum subframe size (equal to one frame), the gain is 0.4 dB, and with the nominal subframe size (equal to one RS block) 128-TCM actually loses 0.5 dB with respect to 64-QAM.

4. CHANNEL IMPAIRMENTS

The three major impairments on the upstream channel in HFC networks are microreflections, narrowband ingress noise, and burst noise [3], [4]. In this section, we analyze the performance of the DOCSIS 2.0 TDMA and S-CDMA in the presence of these impairments.

4.1. Channel Microreflections

A common model for the microreflections on the cable upstream channel (see [3], [4]) is to use an amplitude of –10 dBc up to 0.5 μs, –20 dBc up to 1.0 μs, and –30 dBc beyond 1.0 μs. To simplify the simulations, we used a discrete channel model with 3 echoes, where the first echo is 10 dB below the main signal path and has a delay of one symbol period T, the second echo is 20 dB down and has a delay of 2T, and finally the third echo is 30 dB down and has a delay of 3T. Using this model, we have investigated the impact of channel microreflections on TDMA and S-CDMA, and the results are shown in Fig. 3 for uncoded QPSK. These results show that the difference between the two multiple access techniques is rather small, but the SNR degradation due to intersymbol interference (ISI) is smaller in TDMA even when the number of codes is reduced by 50% in S-CDMA.

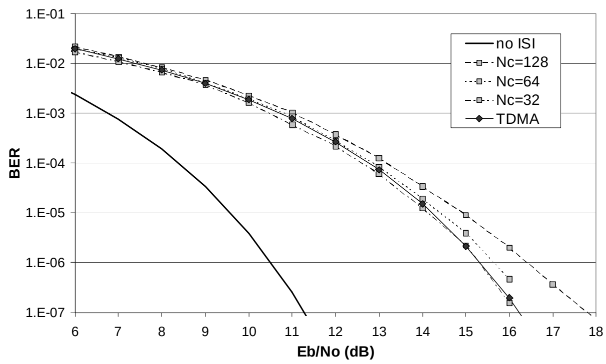

Fig. 3: Impact of channel microreflections on TDMA and S-CDMA.

In the ranging mode, the CMTS identifies the upstream channel impulse response, computes the optimum equalizer coefficients, and sends them to the corresponding CM. But the pre-equalizer setting is never perfect, and there is always some residual ISI at the threshold detector input of the upstream demodulator. Note that any residual ISI is easily handled by the equalizer in a TDMA burst receiver, by identifying the channel impulse response during the preamble. But the individual signals transmitted in parallel by different CMs in S-CDMA having different distortions, there is no way to cancel the residual ISI before the despreader. Furthermore, any processing after the despreader would involve an extremely complex multiuser detector.

4.2. Ingress Noise

Ingress noise originates from man-made sources like shortwave AM and HF radio emissions which leak into the cable and affect the upstream channel. It is modeled as narrowband interference, generally of 20 kHz bandwidth. In a 3.2 or 6.4 MHz wide upstream channel, typically two or three interferers can be encountered, and the carrier-to-interference ratio (CIR) can be as low as 0 dB.

The influence of narrowband interference on code-division multiple access (CDMA) signals is an issue on which there are erroneous ideas in the technical community, because CDMA is often assimilated to direct-sequence spread spectrum (DS-SS) signaling. But as indicated in [5] and [6], the TDMA vs. CDMA issue (whether CDMA uses orthogonal spreading codes or pseudo-noise codes) can not be assimilated to DS-SS vs. non-DS-SS signaling. It is shown in these papers that CDMA and TDMA have the same performance in terms of the CIR at the threshold detector input, and that TDMA actually has better BER performance.

The reason for this surprising result is that while the interference at the threshold detector input has a quasi-uniform amplitude distribution in TDMA, the despreading operation makes it Gaussian in CDMA, and the Gaussian distribution is more sensitive than the uniform distribution in the range of interest. In fact, the most robust multiple access technique to narrowband interference is orthogonal frequency-division multiple access (OFDMA) [7] which can inhibit the carriers that are affected by interference.

In conventional TDMA receivers, compensation of narrowband interference is performed by means of notch filters, and the resulting ISI is compensated using an equalizer. One problem with this approach is it mixes the channel distortion and ingress noise problems, ignoring the fact that the former depends on the transmitting CM while the other does not. A better approach, implemented in the advanced CMTS receiver presented in [8], is described in [9]. It consists of implementing a noise prediction filter to estimate the ingress noise and subtract it from the received signal, and an adaptive equalizer optimized under the zero-forcing criterion to estimate the channel impulse response, and compute and send the pre-equalizer coefficients to the CM of interest. In this receiver structure, the noise prediction filter coefficients are saved from burst to burst, while the pre-equalizer coefficients are recomputed at each burst using the preamble.

Performance of the ingress noise canceller of [9] is illustrated in Fig. 4. This figure shows the influence of three 20-kHz wide interferers of individual power of –15 dBc on 16-QAM, and their compensation using the ingress noise canceller. We can see that the system operates at the BER of 10^{-5} with an SNR degradation limited to 2.1 dB, which is an outstanding performance for this environment.

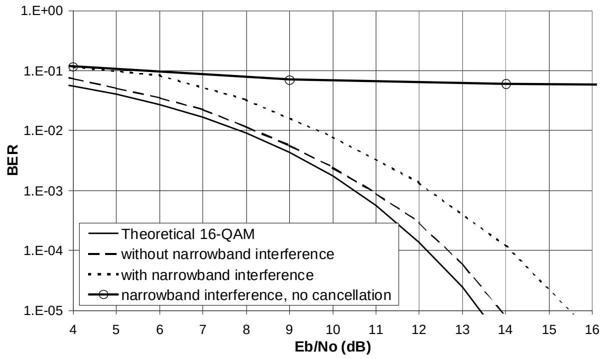

Fig. 4: Effect of narrowband interference on 16-QAM TDMA and its compensation.

In S-CDMA, the ingress noise becomes wideband (and its samples uncorrelated) at the despreader output, and therefore it must be cancelled before the despreading operation. Although this is possible in principle, noise prediction using the receiver decisions in this case involves a significant delay and cannot be made as reliable as in TDMA.

4.3. Burst Noise

Burst noise in cable plants occur mostly at frequen-cies below 10 MHz. We will use a simple model where bursts occur periodically at a predetermined repetition rate. The three parameters in this model are the burst duration, the burst amplitude, and the repetition frequency.

The basic countermeasure against burst noise is error correction coding with interleaving. Such a mechanism is included in both modes of DOCSIS 2.0. In TDMA, noise immunity is determined by the RS code and the byte interleaver. The RS code parameters are the number of information bytes *RSk* and the number of correctable byte errors *RSt*. The block length *N* is given by *N* = *RSk* + *2RSt*. The block interleaver parameters are its width *IntW* and its depth *IntD*. The interleaver width is equal to the RS block length. At the interleaver, the input data is written row by row, and read column by column. The deinterleaver performs the inverse operation.

The bytes affected by a noise burst of *TBurst* seconds will be located in *NC* columns in the deinterleaver. *NC* is related to the symbol rate *R* and the number of bits per symbol *m* by the relation

$$NC = \left\lceil \frac{(TBurst.R + 2) \times m}{8 \times IntD} \right\rceil$$

where $\lceil x \rceil$ denotes the smallest integer larger than or equal to *x*. Since burst noise is generally assumed to have a large power with respect to the useful signal, we can assume that the codeword error rate is 0 if $NC \le RSt$ and 1 if $NC > RSt$.

In S-CDMA, in addition to the RS code and to interleaving, mitigation of burst noise also benefits from signal despreading. But as mentioned earlier, S-CDMA does not use the byte interleaver of TDMA, but instead it uses symbol interleaving over subframes of height (number of spreading codes) *F*. Assuming that each subframe contains one codeword, the bytes affected by burst noise are uni-formly distributed over the codewords of the frame. In the absence of trellis coding, it can be shown that the number of affected bytes per codeword is

$$Nbytes = \left\lceil \left\lceil \frac{TBurst.R + 2}{128} \right\rceil \times \frac{F.m}{8} \right\rceil$$

But note that *Nbytes* may exceed *RSt* in S-CDMA, while still ensuring an acceptable codeword error rate. The latter is given by

$$CwER = \sum_{i=0}^{\max(Nbytes, RSt)} C_{Nbytes}^{i} ByteER^{i} (ByteER)^{(Nbytes-i)}$$

where *ByteER* is the error rate of the bytes affected by burst noise, taking into account the despreading gain on noise power. With optional TCM, calculation of the error rate is very involved, and therefore we resorted to computer simulation. Fig. 5 shows the BER results obtained using 16-QAM with no trellis coding and 32-TCM. The burst noise power in these simulations was 0 dBc and the subframe height was 15.

The results show that the BER is lower with 16-QAM, which exhibits a very sharp drop when the burst duration gets shorter than 256 chips (2×15 S-CDMA symbols per RS codeword). The higher BER of 32-TCM is due to the behavior of the TCM decoder, which actually leads to error events of increased length with respect to that of the noise burst. This suggests that the TCM option should be disabled to operate in the presence of burst noise.

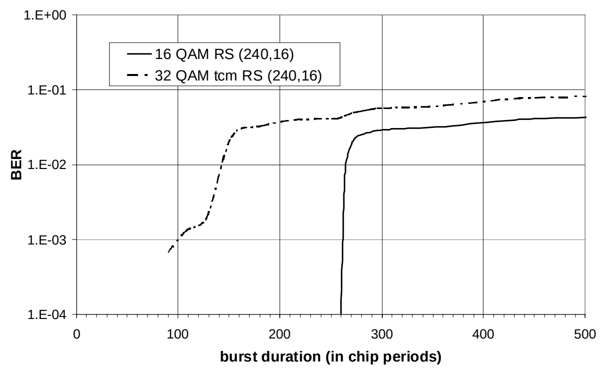

Fig. 5: Performance of TCM with no interleaving in the presence of burst noise.

Next, performance of TDMA and S-CDMA was investigated using the impulse repetition rate vs. impulse duration leading to a BER of 10^{-8} as criterion. The results are depicted in Fig. 6 for 16-QAM and 1518-byte and 64-byte packets. The RS code in these calculations was RS(250, T=16) for the long data packets, and RS(74, T=5) for the short packets. The depth of the byte interleaver used in TDMA was *IntD = 7* for long packets and *IntD = 1* for short packets. Finally, the frame length of S-CDMA was *K = 32*, and the subframe height was *F = M/K*, where *M* is the number of QAM symbols per RS block. Fig. 6 also indicates the minimum and maximum values of the burst repetition rate vs. burst duration that are encountered in cable plants.

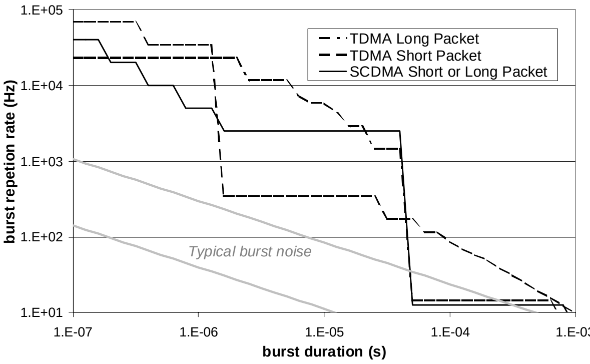

Fig. 6: Performance of TDMA and S-CDMA in the presence of burst noise.

We observe that with exception of a small region near the burst duration of 10^{-4} seconds, both TDMA and S-CDMA can cope with the burst noise of this model. The other observation is that the comparison of TDMA and S-CDMA in this environment is a function of the burst noise characteristics and also of the packets length.

5. CONCLUSIONS

We have analyzed the performance of DOCSIS 2.0 TDMA and S-CDMA on AWGN channels and in the presence of channel impairments, which include microreflections, narrowband ingress noise, and burst noise. Our analysis showed that reducing the number of spreading codes in S-CDMA to increase performance loses too much in data rate to be of practical interest. The set of constellations and RS code parameters included in the standard give the desired flexibility in trading off performance against useful data rate. Also, trellis code option without interleaving gives little improvement on one hand and degrades performance in the presence of burst noise on the other hand. Finally, TDMA and S-CDMA have similar sensitivities to channel impairments, but it is easier to compensate for microreflections and narrowband ingress noise in burst TDMA receivers than in S-CDMA receivers.

REFERENCES

[1] Data-over-Cable Service Interface Specifica-tions – RF Interface Specification, SP-RFI-I05-991105, Cable Television Laboratories, November 1999, Louisville, Colorado.

[2] Data-over-Cable Service Interface Specifica-tions – RF Interface Specification, SP-RFIv2.0-I01-011231, Cable Television Laboratories, Nov. 1999, Louisville, Colorado.

[3] B. Currivan, "Cable Modem Physical Layer Specifications and Design," *in Cable Modems: Current Technologies and Applications*, *John Fijolek & al.* (*Editors*), pages 135-142, IEEE Press, 1999.

[4] T. J. Kolze, "An Approach to Upstream HFC Channel Modeling and Physical-Layer Design," *in Cable Modems: Current Technolo-gies and Applications, John Fijolek & al.* (*Editors*), pages 135-142, IEEE Press, 1999.

[5] H. Sari, F. Vanhaverbeke, and M. Moeneclaey, "Extending the Capacity of Multiple Access Channels," IEEE Communications Magazine, vol. 38, no.1, pp. 74-82, January 2000.

[6] M. Moeneclaey, M. Van Bladel, and H. Sari, "Sensitivity of Multiple Access Techniques to Narrowband Interference," IEEE Transactions on Communications, vol. 49, no. 3, pp. 497-505, March 2001.

[7] H. Sari and G. Karam, "Orthogonal Frequency-Division Multiple Access and its Application to CATV Networks," European Transactions on Telecommunications & Related Technologies (ETT), vol. 9, no. 6, pp. 507-516, November/ December 1998.

[8] F. Buda et al., "Design and Performance of a Fully-Digital DOCSIS CMTS Receiver," 2001 NCTA Technical Papers, pp. 212-220, June 2001, Chicago, Illinois.

[9] A Popper, F. Buda, and H. Sari, "An Advanced Receiver with Interference Cancellation for Broadband Cable Networks," Proc. 2002 International Zurich Seminar on Broadband Communications, pp. 23.1-23.6, February 2002, Zurich, Switzerland.