DESIGN AND PERFORMANCE OF A FULLY-DIGITAL DOCSIS CMTS RECIEVER

F. Buda, E. Lemois, A. Popper, J. Boutros, G. Karam, and H. Sari
Pacific Broadband Communications

ABSTRACT

This paper describes a DOCSIS-compliant cable modem termination system (CMTS) receiver architecture with very advanced features. All receiver functions are implemented digitally, and this feature, together with the advanced signal processing techniques used, leads to an ultra compact and highly scalable CMTS. The receiver architecture adopted makes it possible to implement in a single-chip several upstream burst demodulators along with the corresponding downstream modulators and medium access control (MAC)-layer functions. We report simulation and measurement results confirming the extremely high performance of the described CMTS receiver both in the QPSK and the16-QAM modes.

1. INTRODUCTION

Hybrid fiber/coax (HFC) networks, which were originally used for broadcast TV services, have recently evolved to two-way networks that deliver high-speed Internet access to residential users. The customer premises equipment in this application is referred to as cable modem (CM) and the network side equipment is called cable modem termination system (CMTS). Potential technologies for high-speed Internet access are asymmetric digital subscriber loops (ADSL) over twisted-pair telephone lines, satellite access, broadband fixed wireless access, and HFC network access. The CM technology has taken the lead among all those technologies, and in the US alone there are today millions of households connected to the Internet over HFC networks.

Standardization for digital data services over HFC networks was undertaken in the past by several organizations including the Digital Video Broadcasting (DVB) project, the Digital Audio-Visual Council (DAVIC), and the IEEE 802.14 Group. But the slowness of officially accredited standardization groups incited cable operators in the US to form the Multimedia Cable Network Systems (MCNS) consortium in 1995 and define a standard called Data over Cable System Interface Specification (DOCSIS), which has become the *de facto* industry standard in the US. Now, there is also a European version of this standard called Euro-DOCSIS.

The key success factors to any broadband access technology are performance, equipment size, and cost. In the past, there has been a significant effort to integrate and lower the cost of customer premises equipment (CM's), but little effort has been made to reduce the size and cost of CMTS equipment. This is essentially due to the fact that a single CMTS traditionally serves a large number of subscribers, and the cost per user is typically small. This reasoning, which is common to all point-to-multipoint systems, no longer holds when the number of users per network access point becomes small. This is precisely the situation with HFC networks as fiber nodes shrink in size and get closer to the subscribers, and the number of subscribers per port gets smaller. The purpose of this paper is to describe a fully digital CMTS receiver architecture that leads to a very compact and flexible implementation of the CMTS, while ensuring excellent performance.

The paper is organized as follows: First in the next section, we briefly review the DOCSIS Physical (PHY) Layer specification. Next, in Section 3, we present the receiver architecture and describe in some detail the digital front-end and the digital demodulator functions. Section 4 reports simulated and measured performance results of the receiver in the initial ranging mode and in the traffic mode. Finally, we give our conclusions in Section 5.

2. THE DOCSIS STANDARD

DOCSIS is a set of technical specifications [1], [2], which were structured under the leadership of CableLabs to guarantee multi-vendor interoperability. The DOCSIS RF specification includes the physical-layer, the data link control (DLC) layer that includes the medium access control (MAC) sublayer, as well as a convergence layer with upper network layers. Since the topic of this paper concerns the uplink receiver in the CMTS, our description will be limited to the PHY layer with a particular focus on uplink transmission.

2.1. Downstream Channel

Downstream transmission (from the CMTS to CM's) on cable networks can use a channel in a wide spectrum between 50 and 860 MHz. This spectrum is channelized using 6-MHz channel spacing. The modulation format is quadrature amplitude modulation (QAM) with 64 constellation points (64-QAM) or 256 constellation points (256-QAM) [3]. Channel filtering uses a raised cosine filter equally split between transmitter and receiver. The roll-off factor is $\alpha = 0.18$ for 64-QAM modulation and $\alpha = 0.12$ for 256-QAM modulation. The nominal symbol rate on a 6-MHz channel is 5.056941 Mbaud in 64-QAM mode, and 5.360537 Mbaud in 256-QAM mode.

Error correction coding is based on a concatenated coding scheme with an external Reed-Solomon (RS) code [4], [5], an inner pragmatic trellis code [6] and a convolutional interleaver. The RS code used is an RS(128, 122) which is a 3-symbol error correcting code defined over the Galois field GF(128).

2.2. Upstream Channel

Upstream transmission in the DOCSIS standard uses the 5 - 42 MHz frequency band. This spectrum can accommodate a number of upstream channels of different bandwidths. The channel bandwidth W in the DOCSIS specifications can take the values of 200, 400, 800, 1600, and 3200 kHz. The nominal symbol rates for these channel bandwidth values are 160, 320, 640, 1280, and 2560 Kbaud, respectively. That is, the symbol rate is given by R = 0.8W. The two modulations specified are the simple quaternary phase-shift keying (QPSK) modulation and 16-QAM. For both modulations, channel filtering uses a raised cosine filter (equally split between transmitter and receiver) with roll-off factor $\alpha = 0.25$. The multiple access scheme is a combination of frequency-division multiple access (FDMA) and time-division multiple access (TDMA), i.e., FDMA/TDMA. In this scheme, the CMTS assigns each CM to one channel and allocates time slots to it on that channel.

The upstream is coded using a Reed-Solomon code over GF(256) with a correction capacity of T = 1 to 10 symbols. But there is also an uncoded mode, which corresponds to deactivating this forward error correction (FEC) code. The RS code block length ranges from 18 to 255 bytes, and the number of information bytes per code word ranges from 16 to 253. There are two modes for coding the last block of each burst. The first one is a Fixed Codeword Length which consists of appending by a (0, 0, …, 0) sequence the last block so that all blocks to the RS coder input are of equal length. The other is Shortened Codeword Length in which the last block is not appended and remains of shorter length than the preceding blocks. The latter mode has the advantage of reducing overhead.

DOCSIS specifications are very flexible in the sense that the modulation format and the FEC code can be defined on a burst-by-burst basis. The burst length itself is redefined at each burst.

2.3. Ranging and Traffic Modes

CM's on HFC networks operate in two different modes: The ranging mode during which different parameters are set, and the traffic mode during which useful data is transmitted. The CM enters the ranging mode at connection set-up (initial ranging) in order to perform carrier synchronization, timing clock synchronization, and power control. In the ranging mode, there is a large uncertainty on these parameters and the search for the optimum parameters must be therefore performed over an extended range. In the traffic mode, the CMTS has some *a priori* knowledge of these parameters, and the uncertainty is small. Synchronization problems must therefore be examined in the ranging mode.

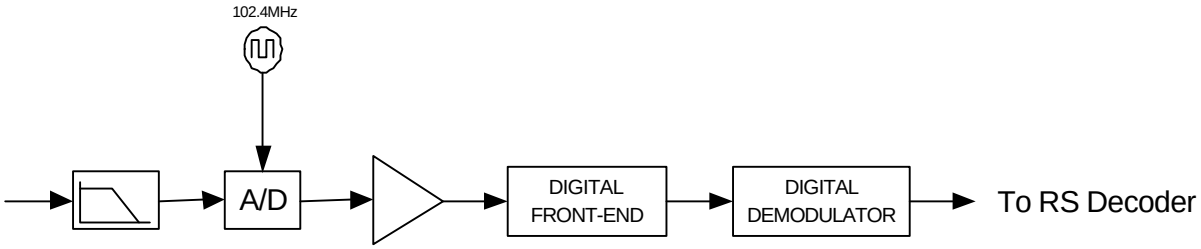

Fig. 1: General block diagram of the CMTS receiver.

Another important function of the CMTS in the ranging mode is to estimate the channel impulse response and compute the optimum equalizer coefficients for that channel. The computed coefficients are then sent to the CM, which is in charge of pre-equalizing the transmitted signal in the traffic mode.

3. RECEIVER ARCHITECTURE

A general block diagram of the receiver is shown in Fig. 1. The received signal is first filtered, amplified, and A/D converted using a clock generated by a free-running oscillator. The nominal frequency of this clock is 102.4 MHz. The variable-gain amplifier used controls the signal power of the entire carrier multiplex. After A/D conversion, the signal is sent to the fully digital front-end which is followed by the digital demodulator.

3.1. Front-End

A functional block diagram of the digital front-end is shown in Fig. 2. The first function of this block is to convert the received digital signal to baseband and generate the in-phase (I) and quadrature (Q) baseband components. This is performed using two multipliers and a numerically controlled oscillator (NCO). The frequency of this oscillator is controlled by the CMTS so as to extract the desired carrier. This signal is then passed to digital filtering and decimation stages, which provide 4 samples per nominal symbol duration. The final stage of the digital front-end is the matched filter, which operates at 4 times the nominal symbol rate and performs square-root raised-cosine Nyquist filtering.

3.2. Digital Demodulator

The front-end is followed by the digital demodulator whose basic function is to perform timing and carrier synchronizations, channel equalization, ingress noise cancellation, and make symbol decisions. A functional block diagram of the demodulator is depicted in Fig. 3.

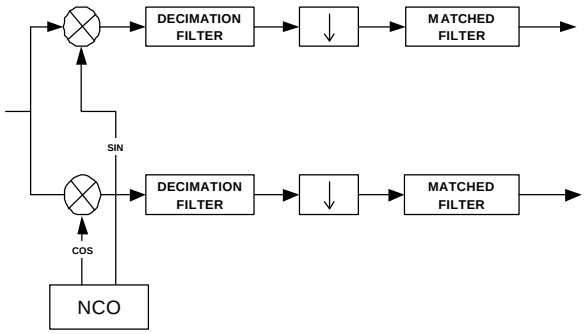

Fig. 2: Block diagram of the digital front-end.

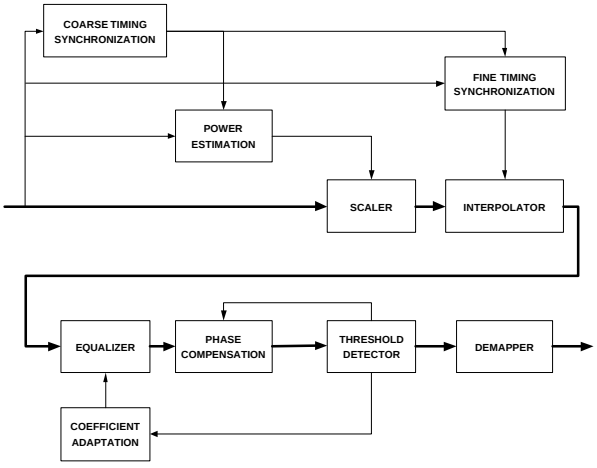

Fig. 3: Block diagram of the digital demodulator.

3.2.1. Coarse Timing Estimation

First, a coarse timing function detects the beginning of each burst with the required precision (typically a precision of half a symbol period). The conventional approach to coarse timing estimation is based on power estimation. The principle is simple: The signal received in the absence of bursts is due to noise, and its value is small compared to the signal received during bursts. Therefore, a power estimation circuit followed by a threshold comparator gives an indication of the start of bursts. The power estimation circuit is composed of two elementary functions: The first one is a squaring circuit which gives the instantaneous signal power, and the second is a low-pass filter which performs short-term averaging.

The first problem associated to this concept is that the precision of the burst start estimate is a function of the filter used. A short filter memory is required to improve precision, but then the estimator becomes very sensitive to additive noise. That is, robustness to noise and precision of the estimator are two contradictory requirements in this technique. The second problem is that the threshold is a function of the received signal power level, which is undesirable. A high threshold leads to the risk of missing bursts and includes an estimation delay. A low threshold reduces the delay, but creates the risk of declaring that a burst is present when no burst is actually transmitted. (This occurs when the noise power exceeds the threshold level.)

To avoid these problems, we developed a new coarse timing detector that involves a correlator and the computation of a *contrast function* that is independent of the received signal power level. The correlator correlates the incoming signal with the preamble sequence stored in the receiver. (The preamble must have good correlation properties, i.e., a very narrow correlation peak and very low correlation values around that peak.)

With a *contrast function* that is independent of the received signal power level, a fixed threshold can be used (without any performance penalty) to detect the correlation peak and the burst start. The threshold comparator in the block diagram determines a short time-window in which the correlation maximum is to be searched. In the traffic mode, the CMTS has some *a priori* knowledge of the burst position and knows the time window over which the *contrast function* needs to be maximized.

3.2.2 Subsequent Demodulator Functions

Next, a fine timing function determines the right sampling instant and passes this information to an interpolator that generates symbol-spaced signal samples. A scaler that precedes the interpolator sets the power of the over-sampled signal to a predetermined value. The scaler is controlled by a power estimation function that is activated during signal bursts. The symbol-spaced signal samples generated by the interpolator are passed to subsequent receiver stages, which include an adaptive equalizer, an ingress noise canceller, and a carrier phase recovery circuit.

The equalizer is a linear equalizer whose coefficient values are computed using the zero-forcing (ZF) criterion [3]. This criterion is more appropriate in the present case than the more popular minimum mean-square error (MMSE) criterion, due to the requirement to send the coefficient values to the CM to implement a pre-equalizer. The reason is that the MMSE equalizer makes a trade-off between channel distortion and additive noise, and therefore the computed coefficients do not perfectly invert the channel transfer function. And contrary to an equalizer at the receiver, a pre-equalizer does not amplify the additive noise. Therefore, the best coefficient setting for the pre-equalizer is that which perfectly inverts the channel.

As is well known, ingress noise represents one of the major disturbances that affect upstream data transmission in HFC networks. Ingress noise is essentially due to local AM radio signals and other types of disturbances that leak into the cable. It is modeled as narrowband interference that may be on or off and essentially constant over a period of time which can be in excess of several minutes. Another characteristic of ingress noise is that, contrary to channel distortion, which is specific to each CM, it is common to all CM's sharing the same upstream carrier. The reason is that the CMTS receives the

sum of all noises that leak into the cable at all customer premises that it serves, and the resulting noise equally affects all time slots no matter where they originate from.

For reliable data transmission on the upstream channel, the receiver must include an efficient ingress noise canceller, particularly for 16-QAM and higher-level modulations. One way to suppress ingress noise is to use a notch filter at the ingress noise frequency, but notch filtering also distorts the useful signal and creates intersymbol interference (ISI), which is undesirable. An alternative approach consists of estimating ingress noise by means of a prediction filter and subtracting this estimate from the received signal prior to threshold detection. The latter approach, which leads to significantly better performance, was adopted in our receiver design.

The final function before the threshold detector (which makes the symbol decisions) is the carrier synchronization function. This includes a frequency estimator that estimates the frequency offset between the CM and the CMTS and a phase recovery circuit that synchronizes the carrier phase of the incoming signal. The estimated frequency offset is used to derive a control signal that is sent to the CM to synchronize its oscillator frequency with that of the CMTS. The decision-feedback frequency estimator used in our design is based on a newly developed algorithm that is very robust against symbol decision errors. Finally, the phase recovery circuit compensates for residual synchronization errors between the CM and the CMTS.

4. PERFORMANCE RESULTS

Performance of the designed CMTS receiver was evaluated using extensive computer simulations and laboratory measurements. In this section, we will report results that assess the performance of the digital front-end and of individual demodulator functions, as well as results that assess overall receiver performance.

4.1. Front-End Performance

Performance of the CMTS receiver was tested for different symbol rates and different loads of the upstream spectrum. The most unfavorable condition for the digital front-end occurs when the desired signal has the lowest symbol rate (160 kbaud) and arrives at the receiver with the minimum signal level, while adjacent carriers arrive with the maximum signal level allowed in DOCSIS specifications.

To evaluate worst-case performance, we have simulated a carrier multiplex where a 160 kbaud desired signal is received together with 8 adjacent carriers each having a data rate of 2560 kbit/s and a power spectral density (psd) that is 12 dB above that of the desired signal. That is, the desired signal power was 24 dB below that of each one of the other carriers. Our simulations have indicated that the receiver performance in these conditions is essentially the same as in the case of an isolated carrier over the upstream channel. The measurement results using a lab prototype were very much in agreement with the simulation results.

4.2. Overall Performance in the Traffic Mode

Next, we simulated the overall bit error rate performance (BER) of the receiver in the traffic mode and plotted it as a function of the signal-to-noise ratio (SNR). The results take into account the total imperfections of the receiver including those of the front-end and of the synchronization functions. The BER vs. E_b/N_0 (transmitted energy per bit to the noise spectral density ratio) curves are given in Fig.4 for QPSK and in Fig. 5 for 16-QAM. The figures also show the theoretical BER curves which correspond to the performance of an ideal modem.

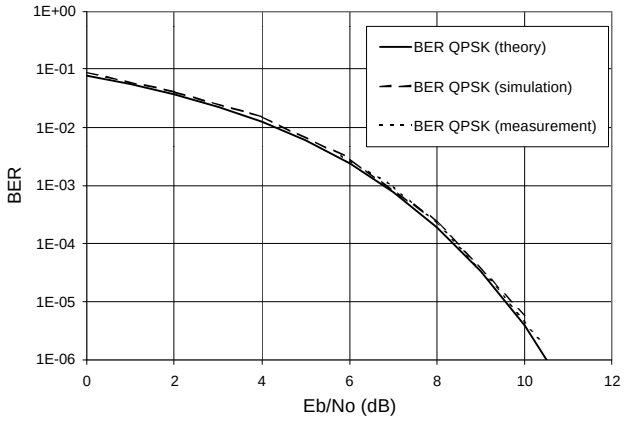

Fig. 4: Overall BER performance of the CMTS receiver in the QPSK mode.

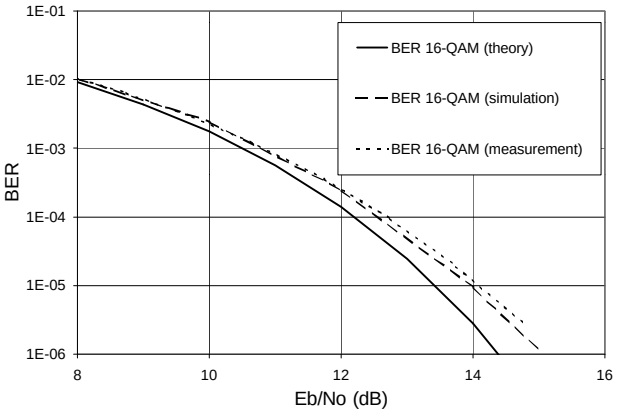

Fig. 5: Overall BER performance of the CMTS receiver in the 16-QAM mode.

These figures show that the overall degradation at the BER of 10^{-6} is limited to 0.2 dB in QPSK and 0.6 dB in 16-QAM. These results are obtained in the absence of error-correction coding. This means that the 0.2 dB SNR degradation in QPSK and 0.6 dB degradation in 16-QAM will hold for BER values as low as 10^{-10} or 10^{-12} after RS decoding.

Figs. 4 and 5 also give the measurement results. We can see that the measurement results coincide with the simulation results in QPSK, and that the difference between simulated and measured results is limited to 0.1 dB in 16-QAM.

4.3. Burst Detection in the Ranging Mode

The most significant performance indicator in the ranging mode is the time needed by a CM to register with the network. The registration time must be evaluated in two extreme cases: The worst case, which corresponds to all CM's using the ranging opportunities, a situation which typically occurs after a CMTS Reset, and the best case, which corresponds to only one CM using the ranging opportunity. The latter case is in fact sufficient to determine the performance of the CMTS receiver.

The important parameters to consider here are the elementary probabilities Pnd_{el} and Pfa_{el} which respectively correspond to missing a burst and to a false alarm at a given time *t*. Missing a burst occurs when the burst is actually transmitted and the *contrast function* used for burst detection takes a value lower that the decision threshold *S*. A false alarm corresponds to the contrast function taking a value that exceeds the decision threshold while no burst is actually transmitted. Note that Pfa_{el} is independent of the SNR, because in the absence of useful signal, both the numerator and the denominator of the contrast function C(t) are proportional to the noise variance, and therefore the noise variance cancels out. In contrast, Pnd_{el} is a function of the SNR. To reduce Pnd_{el}, we need to decrease the threshold *S*, and to reduce Pfa_{el}, we need to increase *S*. That is, reducing the elementary non-detection probability is a contradictory requirement with reducing the elementary false alarm probability. This is shown in Fig. 6 where we have plotted Pnd_{el} for E_b/N_0 = 10 dB and E_b/N_0 = 8 dB.

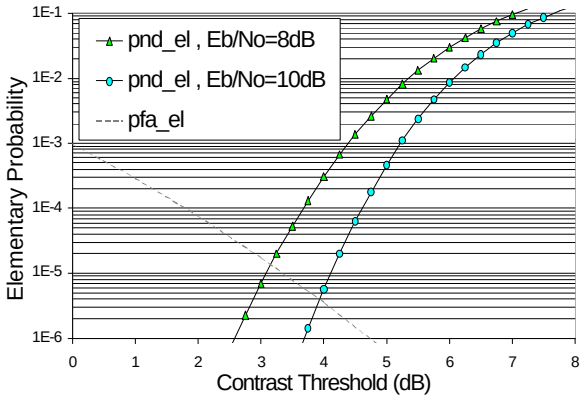

Fig. 6: Elementary false alarm and non-detection probabilities.

But Fig. 6 alone is not sufficient to determine the optimum value of the threshold *S*. To determine this value, we need to consider the full probability of missing a transmitted burst, which we denote *Pnd*. To evaluate this probability, we need to consider the following two situations:

- The transmitted burst is not detected due the *contrast function* taking a value lower than the threshold. This occurs with a probability of Pnd_{el}.

- A false alarm occurs during one of the *N* symbols preceding the burst start, where *N* is the number of symbols in the ranging burst. A false alarm will activate the demodulator and deactivate the *contrast function* calculations for the following *N* symbols, and therefore a true

burst start during N symbols after a false alarm will not be detected. The probability of missing a burst due to false alarms is therefore $N.Pfa_{el}$.

Taking into account, these two types of missing a burst, the total probability of missing a burst is given by

$$Pnd = Pnd_{el} + N.Pfa_{el}.$$

According to DOCSIS specifications, a ranging burst must carry at least 34 bytes, which map on 136 QPSK symbols. Taking into account this minimum value as well as the preamble and the redundancy for error-correction coding, the number of symbols in the ranging burst is close to 200. Therefore, $Pnd = Pnd_{el} + 200Pfa_{el}$ is a good approximation for the probability of missing a transmitted burst, and the threshold S must be set so as to minimize this probability. In Fig. 7, we have plotted the probability of missing a burst as a function of the threshold value, for E_b/N_0 = 10 dB and E_b/N_0 = 8 dB.

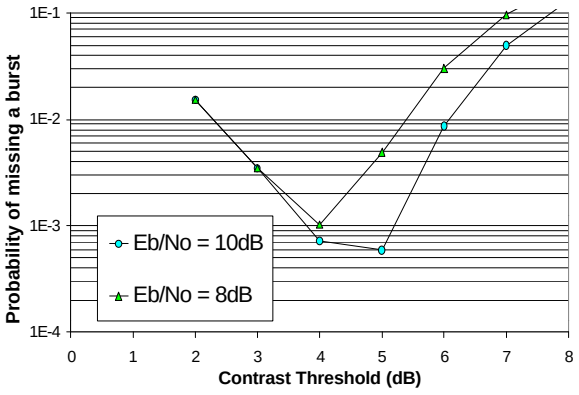

Fig. 7: *Pnd* as a function of the threshold.

This figure indicates that any value of the threshold between 3 and 4 will give a P_{nd} lower than 10^{-2} for E_b/N_0 values higher than 10 dB. A P_{nd} lower than 10^{-2} means that the CM will use the first ranging opportunity 99% of the time, and will need at least two ranging opportunities with a probability of 10^{-2}. Continuing further, the CM will need 3 ranging opportunities only with a probability of 10^{-4}. Consequently, the average registration time of a CM (assuming that only one CM is trying to register at a time) is well approximated by 1.01 times the interval between two consecutive ranging opportunities.

4.4. Pre-equalizer Performance

As mentioned earlier, the pre-equalizer coefficients are computed by the CMTS receiver and sent to the CM to pre-equalize the transmitted signal in the traffic mode. Here, we give the performance corresponding to two different options. In the first, the coefficients are estimated using a single ranging burst, and are immediately sent to the transmitter after that burst. In the second, two ranging burst are used in order to obtain a better estimate of the optimum pre-equalizer coefficients.

To assess the BER performance of the pre-equalizer, we used a channel model with 3 echoes: The first echo is 10 dB below the main signal path and has a delay of 1 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. Note that the performance results obtained with this model are independent of the symbol rate. Fig. 8 shows the BER curves corresponding to 16-QAM and an 8-tap pre-equalizer. The figure also shows the theoretical 16-QAM curve corresponding to an ideal modem operating on a channel with no distortion.

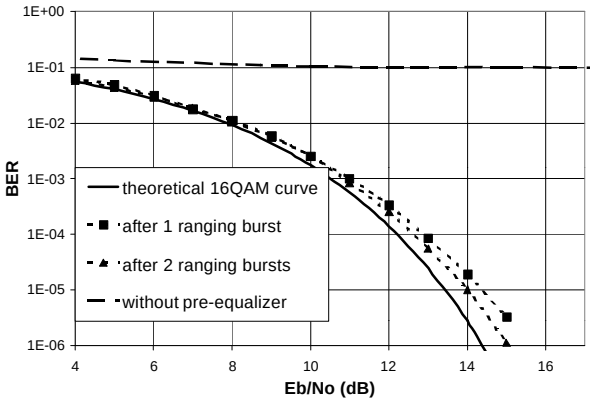

Fig. 8: *BER* performance of an 8-tap pre-equalizer with 16-QAM modulation.

Notice that without a pre-equalizer, the system has an irreducible BER on the order of 10^{-1}. The results show that the SNR degradation at the BER of 10^{-5} is approximately 1 dB when only a single

ranging burst is used to optimize the pre-equalizer coefficients. This degradation is reduced to 0.6 dB when two ranging bursts are used for coefficient optimization.

Next, we investigated the pre-equalizer signature using the 1-echo channel model which appears in DOCSIS specifications. The signature gives a plot of the echo amplitude (in the dB scale) vs. echo delay (normalized by the symbol period T) that leads to an SNR degradation of 0.8 dB at the BER of 10^{-5}. The results corresponding to 16-QAM and an 8-tap pre-equalizer are depicted in Fig. 9.

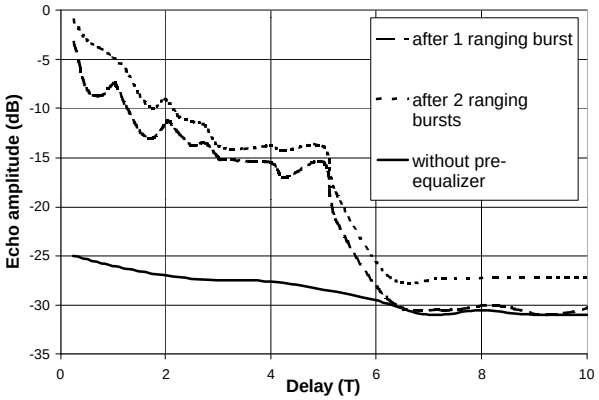

Fig. 9: Signature of an 8-tap pre-equalizer with 16-QAM modulation.

The figure shows that non-equalized 16-QAM system only tolerates an echo of –25 dB (the normalization is with respect to the main signal path). It also shows that for delays up to 5 symbol periods, an 8-tap pre-equalizer can cope with an echo amplitude of –15 dB. The sharp drop of performance is due to pre-equalizer size and reference-tap position used in these simulations. Another interesting observation that can be made here is that using two ranging bursts to optimize the pre-equalizer improves the signature by 1 to 3 dB.

4.5. Ingress Noise Canceller Performance

Performance of the implemented ingress noise canceller was simulated using a 16-QAM signal at the symbol rate of 2.56 Mbaud. The ingress noise model used includes three interferers respectively centered at a distance of 60 kHz, 500 kHz, and 1000 kHz, from the carrier frequency. Each interferer has a width of 20 kHz and an individual power that is 15 dB below the useful signal power.

The results obtained using this model are depicted in Fig. 10. Notice that a 16-QAM modem cannot operate in the presence of this type of ingress noise without an interference canceller. Clearly, the BER curve shows a floor close to 10^{-1}, which is entirely unacceptable. With a simple canceller based on noise prediction, the BER curve becomes parallel to the ideal 16-QAM curve, and the SNR degradation becomes less than 2 dB. Better performance can be achieved by increasing the number of taps of the noise prediction filter.

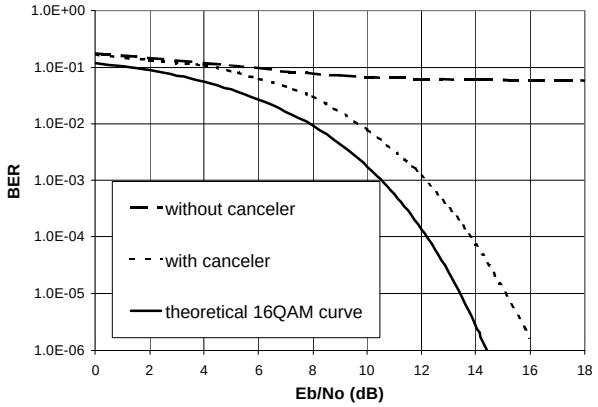

Fig. 10: Influence of ingress noise on 16-QAM and its compensation using noise prediction.

5. SUMMARY AND CONCLUSIONS

We have presented a fully digital receiver architecture, which substantially reduces the size and cost of the CMTS while ensuring excellent overall performance even under worst-case conditions. First, the simulation and measurement results confirmed that the multi-channel front-end gives quasi-ideal performance in the very unfavorable condition where in addition to the useful signal, the cable carries a multiplex of 8 modulated carriers with an individual power that is 24 dB above that of the useful signal. Next, the total SNR degradation of the receiver at the (uncoded) BER of 10^{-6} was found to be limited to 0.2 dB in the QPSK mode and 0.6 dB in the 16-QAM mode. It was also found that with the synchronization algorithms implemented, the average time needed by a CM to register with the network (assuming only one CM is

trying at a time) is only 1.01 times the ranging opportunity, which is an extremely short registration time. Finally, the receiver also includes an efficient adaptive equalizer to compensate for channel distortion and an ingress noise canceller.

The presented receiver was extensively tested using an FPGA implementation and integrated in an ASIC using high-speed 0.18 μ CMOS technology. Prototypes of this chip, which also includes other physical-layer functions as well as MAC-layer functions, are due from foundry in June.

REFERENCES

[1] Data-over-Cable Service Interface Specifications – Radio Frequency Interface Specification, SP-RFI-105-991105, Cable Television Laboratories, November 1999, Louisville, Colorado.

[2] Digital Multiprogram Systems for Television Sound and Data Services for Cable Distribution, ITU-T Recommendation J.83 (Annex B), April 1997, ITU, Geneva, Switzerland.

[3] J. Proakis, “Digital Communications”, Third Edition, McGraw Hill, 1995.

[4] R. Blahut, “Error Control Codes,” Addison Wesley, Reading, MA, 1984.

[5] G. C. Clark, Jr., and J. B. Cain, “Error Correction Coding for Digital Communications,” Plenum Press, New York, 1981.

[6] J. K. Wolf and E. Zehavi, “Pragmatic Trellis Codes Utilizing Punctured Convolutional Codes”, IEEE Communications Magazine, no. 32, pp. 94 – 99, February 1995.