

Synchronized Discrete Multi-Tone (SDMT) Modulation for Cable Modems: Making the Most of the Scarce Reverse Channel Bandwidth

Krista S. Jacobsen

Amati Communications Corporation

2043 Samaritan Drive

San Jose, CA 95124

Phone: (408) 879-2039 FAX: (408) 879-2912

jacobsen@amati.com

Abstract

This paper describes Synchronized Discrete Multi-Tone (SDMT), a technique that enables the use of discrete multi-tone (DMT) in the reverse channel of cable or hybrid fiber-coax (HFC) networks. It is shown that SDMT achieves high spectral efficiencies, on the order of 4-5 bits/s/Hz, and offers resilience to common reverse channel impairments, including impulse noise, radio-frequency ingress, and spectral variability. The paper concludes that SDMT is a cost-effective solution that enables efficient, robust, high-speed upstream communications on cable and HFC networks.

1 Introduction

Both the MCNS data-over-cable specification and the upcoming IEEE 802.14 cable modem standard require use of QPSK for upstream transmission. Although QPSK is robust, inexpensive, and well-understood, it is also inefficient, supporting spectral efficiencies of less than 2 bits/s/Hz, which limits both the number of subscribers that can be supported in the shared reverse channel bandwidth and the maximum burst rate at which those subscribers can transmit. For this reason, the 802.14 Working Group has established a Study Group whose charge is to investigate and propose for standardization a more efficient, higher-performance modulation scheme for upstream transmission. This paper describes one such technique called Synchronized Discrete Multi-Tone (SDMT).

SDMT is a powerful technique that enables a number of DMT-based remote terminals to communicate with a single central-site receiver. SDMT maintains all the advantages that DMT provides in point-to-point environments. These advantages include the ability to support high data rates on channels with a variety of transmission impairments such as passband ripple, spectral nulls, impulse noise, and radio-frequency interference, all of which are expected in the reverse channel of cable and




Figure 1: A channel partitioned into subchannels

HFC networks.

Section 2 reviews DMT and its advantages. SDMT is described in Section 3. An example illustrates why synchronization is necessary in multipoint-to-point environments, and a method to ensure synchronization of multiple DMT-based remote terminals is presented. In Section 4, an example SDMT configuration for cable or HFC networks is described. Section 5 concludes that SDMT enables high-performance, cost-effective multipoint-to-point communications on cable and HFC networks.

2 Discrete Multi-Tone (DMT) Modulation

This section briefly reviews Discrete Multi-Tone modulation, which is the basis of SDMT. For additional details, the reader is referred to [1] through [7].

Discrete Multi-Tone (DMT) modulation uses an *inverse discrete Fourier transform (IDFT)* to partition a transmission channel into a set of orthogonal, equal-bandwidth subchannels so that the frequency response of each subchannel is roughly constant across its bandwidth, as illustrated in Figure 1. The resulting subchannels are almost memoryless, and the small amount of intersymbol and intersubchannel interference caused by the channel's non-unity impulse response length can be eliminated by use of a *cyclic prefix*. As its name implies, the cyclic prefix precedes each symbol. It is a copy of the

last ν time-domain (that is, post-IDFT) samples of each DMT symbol. Hence, the cyclic prefix carries redundant information and is transmission overhead. If there are no more than $\nu + 1$ non-zero samples in the channel impulse response, then intersymbol interference (ISI) caused by each symbol is confined to the cyclic prefix of the following symbol. By discarding the cyclic prefix samples in the receiver, ISI is eliminated completely. Therefore, a traditional equalizer is not required in the DMT receiver. If the cyclic prefix is not long enough, the received signal incurs some distortion.

In general, the channel can be partitioned into any convenient number of subchannels, as long as the condition of nearly-flat frequency responses over the subchannel bandwidths is met. However, one cost-effective implementation of DMT uses the fast Fourier transform (FFT) to compute efficiently the DFT and IDFT required in the receiver and transmitter, respectively. To exploit the computational efficiency offered by the FFT, the number of subchannels should be an integer power of two. Cost constraints and performance requirements then dictate the precise number of subchannels into which a channel is partitioned.

In the DMT transmitter, a bitstream is encoded as a set of quadrature amplitude modulated (QAM) subsymbols, where each QAM subsymbol represents a number of bits determined by the signal-to-noise ratio (SNR) of its associated subchannel. During each DMT symbol period, a set of subsymbols is input as a block to a complex-to-real IDFT. The cyclic prefix is prepended to the output, and the result is converted from digital to analog format. If passband transmission is required, the resulting signal, a sum of real time-domain sinusoids, is then upconverted to the desired carrier frequency and transmitted over the channel. In the receiver, after downconversion and analog-to-digital conversion, the cyclic prefix is stripped, and the samples are input to a DFT. Each output value is then scaled by a single complex number to compensate for the magnitude and phase of each subchannel's frequency response, and a memoryless detector decodes the resulting symbols. The set of complex multipliers is known as the *frequency-domain equalizer (FEQ)*. Changes in the channel magnitude or phase are accommodated by updating the FEQ taps as the system operates. A generalized passband DMT transmitter and receiver pair is shown in Figure 2.

2.1 Advantages of DMT

DMT provides a number of advantages for transmission on difficult channels such as the reverse channel of cable or HFC networks. First, DMT can maximize the bit rate supported on any channel. In contrast to single-carrier techniques, which support a fixed number of bits over a set frequency band, the number of bits supported by DMT varies with frequency. Subchannels are assigned as many bits as they can support reliably. As a result, the aggregate bit rate on the channel is generally higher

than it would be using a single-carrier technique. For example, QPSK transmits 2 bits per symbol. Taking into account the required excess bandwidth, the spectral efficiency of QPSK is less than 2 bits/s/Hz, irrespective of the channel capacity. In contrast, the spectral efficiency of DMT is closely tied to the channel capacity: more bits per symbol are supported on channels that have higher capacities. It is expected that the SNR in some spectral regions of the reverse channel will be high enough to support 5 or even 6 bits. DMT takes advantage of these high SNR regions by assigning the appropriate number of bits to those regions, thereby maximizing the channel bit rate. DMT also automatically avoids those regions of the spectrum in which meaningful communications cannot be supported. During initialization, the bit distribution is computed based on the observed channel and noise conditions. Those subchannels with SNRs insufficient to support any bits are simply turned off. By tailoring the bit distribution to the channel characteristics, DMT optimizes the transmit bandwidth.

During steady-state operation, the bit distribution is adapted as necessary to maintain a desired performance. The subchannel SNRs are monitored in a data-driven manner by the receiver. Upon detecting a degradation in a subchannel SNR, the receiver computes a modified bit distribution that better achieves the desired error performance. Depending on the SNR of a degraded subchannel, some or all of its bits may be moved via a *bit swap* algorithm to one or more other subchannels that can support additional bits. The bit distribution change is reported to the transmitter, which then notifies the receiver when the new bit distribution will be activated.

The adaptivity afforded by bit swapping enables DMT to alleviate problems caused by changes in the channel frequency response or noise while a system operates. Even unpredictable, bursty, high-power noise such as ingress from amateur radio or other radio-frequency interferers can be tolerated by a DMT system, assuming that either the resulting channel can continue to support the required data rate or the data rate can be decreased and that the interference does not overload the receiver and cause nonlinear operation. Dramatic changes in the channel itself can also be accommodated by an adaptive DMT system, provided the adaptation speed of the bit swap algorithm is sufficient to track the changes. Note that changes in the channel frequency response or noise during operation that actually improve any of the subchannel SNRs can be used to increase the system's overall data rate, if desired. Alternatively, if a fixed data rate is desired, improvements in subchannel SNRs yield higher noise immunity by increasing the noise margin.

Practical single-carrier transceivers, such as QPSK, are much less proficient than DMT at optimizing the transmission bandwidth and operating in the presence of radio-frequency ingress: most finite-length equalizers cannot produce required nulls, notches, and "dead-zones" in the frequency spectrum as accurately and eas-




Figure 2: Passband DMT block diagram

ily as DMT can simply by turning off subchannels.

Another advantage of DMT is its immunity to impulse noise, which is generally better than that of single-carrier systems. Because the DMT receiver processes received samples in blocks, after the receiver transform the energy of any impulse noise that corrupted some of the received samples is spread evenly among the subchannels. Hence, impulse noise causes a temporary decrease in each subchannel SNR. If the subchannel noise margins are large enough, impulse noise may not affect the system error performance. If, however, the expected impulse noise is severe enough to exceed the noise margin, then forward error correction with interleaving can be applied to mitigate errors. [8]

3 Synchronized DMT (SDMT)

Section 2 described point-to-point DMT, in which a single transmitter communicates with a single receiver over a dedicated channel. In a point-to-point environment, because the DMT transmitter modulates data using an IDFT with a cyclic prefix, the resulting subchannels are independent and orthogonal. However, the reverse channel of an HFC network is shared by multiple remote transmitters. Hence, precautions must be taken to use DMT in multipoint-to-point environments. Otherwise, the elegant subchannel orthogonality achieved so easily in point-to-point DMT can be lost.

To motivate Synchronized DMT (SDMT), consider the following example: Assume that exactly three remote terminals (A, B, and C), located on an HFC network such that their signal propagation times to the headend differ, wish to transmit data to a single headend receiver. To begin, assume that the three remote units have been permanently assigned disjoint sets of subchannels for their upstream transmissions¹, as shown in Figure 3. By restricting the remote terminals to transmit on

¹Such a restrictive allocation of subchannels for data transmission is undesirable for a number of reasons that are outside the scope of this document. However, the FDMA-like configuration helps to demonstrate the need for remote terminal synchronization, and the argument for synchronization holds when less restrictive subchannel allocations are allowed.

disjoint sets of subchannels, the time-varying nature of the multi-access problem is eliminated for the time being, and synchronization requirements related only to the modulation can be examined. Under this assumption, Figure 3 illustrates conceptually the signal that might be received at the headend if Remotes A, B, and C all transmit data for a time equal to one symbol period on their respective subchannel sets beginning at time $t = 0$. Because Remote C is closest to the headend, its signal (information on subchannels marked "C" in the figure) arrives first at the receiver. The signal from Remote B arrives next, and data on subchannels used by Remote A arrives last. Because the remote units "own" their respective subchannel sets, there is no possibility that more than one remote unit transmits at a time on any subchannel. The received time-domain signal is then the sum of the signals transmitted by the individual remote terminals, which by construction contain only frequency components corresponding to the subchannels assigned to each remote unit. As a result, subchannel orthogonality is ensured in this case. Hence, in theory, the remote terminals can transmit as desired on their designated subchannel sets, and the headend receiver can demodulate and reconstruct correctly the original signal provided its sampling rate is at least twice the bandwidth of the transmitted signals. Thus, when disjoint sets of subchannels are assigned permanently to the remote terminals, synchronization of the remote terminals is not required; however, it dramatically simplifies the complexity of the headend receiver. From the discussion in Section 2, the DMT receiver is a block processor: it collects a number of time-domain signal samples and manipulates these samples to reconstruct the set of QAM subsymbols that was transmitted during a symbol period. Hence, the DMT receiver must define where the collection of each block of samples begins and ends. Ideally, the receiver begins collecting data samples immediately after the cyclic prefix samples, which are, in general, corrupted by interference from the preceding symbol, as described in Section 2. In a multipoint-to-point configuration in which multiple remote terminals transmit on permanently-assigned disjoint subchannel sets, however, if transmissions from the remote units




Figure 3: Example of multipoint-to-point DMT communications without synchronization




Figure 4: Loss of well-defined cyclic prefix when unsynchronized remote terminal transmissions are allowed in multipoint-to-point configurations

are not synchronized to arrive at the headend at the appropriate times, then the cyclic prefix samples corresponding to the various subchannel sets most likely do not coincide in time in the received signal; hence, “the cyclic prefix” no longer exists as a distinct, well-defined set of received samples that can be discarded in the receiver. Instead, cyclic prefix samples transmitted by the various remote terminals may be scattered throughout the received data. Figure 4 illustrates the scattering of the cyclic prefix assuming the three-remote configuration shown in Figure 3. Because the received signal is the sum of unsynchronized remote terminal transmissions, the individual cyclic prefixes cannot be removed from the received signal. Hence, use of a cyclic prefix does not help to eliminate ISI if the remote terminal transmissions are not synchronized. Instead, to mitigate ISI a bank of parallel multi-tap equalizers (one equalizer per subchannel) must be implemented in the receiver, which dramatically increases the complexity of the headend modem. Moreover, equalization is a suboptimal detection technique, so the performance of an unsynchronized equalized system may be degraded with respect to a synchronized system that preserves the cyclic prefix.

The preceding discussion presumed that disjoint subchannel sets are permanently assigned to the remote terminals, implying that no additional channel access protocol is required: remote terminals simply transmit as desired on their respective subchannel sets, and no bandwidth conflicts occur. However, cable and HFC networks almost certainly require a less restrictive allocation of the

available bandwidth, particularly because remote terminal transmission requirements vary significantly with time. Hence, a more sophisticated channel access protocol must be used to coordinate use of the subchannels. In this case, synchronization of remote terminal transmissions becomes even more critical. Because the remote terminal distances from the headend differ, ensuring for all time that no more than one remote terminal uses any particular subchannel is extremely difficult without synchronization. In addition, the cyclic prefix problem described previously also exists when a less restrictive allocation of subchannels is allowed with unsynchronized transmissions.

If remote terminal transmissions are synchronized so that the symbol boundaries of signals arriving at the headend receiver coincide, however, then the key properties of point-to-point DMT also hold for multipoint-to-point transmission. This observation led Cioffi to propose Synchronized DMT (SDMT). [9] The general idea of SDMT is to coordinate remote terminal transmissions so that the signal arriving at a single receiver appears to have been transmitted by a single remote terminal. The receiver then proceeds, for the most part, as if it were in a point-to-point environment. To ensure that symbols transmitted by the various remote terminals on an HFC network arrive at the headend with synchronized symbol boundaries, remote units must delay their transmissions a number of samples inversely proportional to their physical distances (in terms of length of coaxial cable) from the headend. As a result of the synchronization,




Figure 5: Preservation of the cyclic prefix in multipoint-to-point environments by delaying remote terminal transmissions

the cyclic prefix as defined for point-to-point transmission is preserved; hence, ISI can be eliminated just as in the point-to-point case by discarding the cyclic prefix in the receiver, and no equalizer other than the single-tap FEQ is required. Figure 5 illustrates this concept for the three-remote configuration considered previously. Furthermore, if more than one remote terminal transmits during a particular symbol period using disjoint sets of subchannels, then the headend receiver can decode the received signal by applying the appropriate mixture of FEQ taps, assuming the headend controller knows which remote transmitted on each subchannel.

Synchronization of remote terminal transmissions to enable SDMT can be achieved by a two-step procedure. Initially, remote terminals must synchronize their local clocks to a master clock so that all remote units transmit with respect to a common reference. The master clock must be broadcast by the headend controller in a downstream control channel: because of the limitations imposed by the HFC configuration, broadcasting the clock in a downstream channel is the only means by which all remote terminals can access a common clock signal. In addition, each remote terminal must acquire a sample delay, computed by the headend controller when the terminal is installed on the network, before it can begin operating. A remote terminal's required sample delay is the difference between the maximum round-trip signal delay to and from the headend of all remote units on the network and the remote terminal's round-trip signal delay.

4 Example System for Upstream HFC Transmission

This section presents an example SDMT configuration to support reverse channel communications. Assuming a sub-split HFC configuration, the entire HFC reverse channel spans the frequency band from 5 MHz to 42 MHz. However, designing remote terminals that use the entire 37-MHz band is impractical for a number of reasons. First, the computational complexity of a transceiver operating in 35+ MHz bandwidth is quite high. Second, it is unlikely that the data rate requirements of the remote terminals warrant use of the entire reverse channel at any time. Based on discussions within 802.14, the maximum expected instantaneous bit rate of the services considered is 10 Mbps. If a spectral efficiency of just 2 bits/s/Hz can be achieved in some part of the reverse channel, then a maximum of 5 MHz bandwidth is required to support 10 Mbps (neglecting any excess bandwidth). Hence, none of the remote terminals is ever likely to need access to the entire reverse channel. A more practical solution than designing wide-band transceivers is to segment the reverse channel bandwidth into smaller bands, each of which supports an independent SDMT signal.

Assuming that minimizing the cost of the remote terminals is a design goal, dividing the entire reverse channel into 1024 34.5-kHz-wide subchannels appears to be a good compromise between performance and complexity. To ensure the complexity of the remote terminals is not prohibitive, the set of 1024 subchannels can be

Table 1: Example system parameters

Total number of subchannels	1024
Number of subchannels per group	128
Maximum nominal remote terminal FFT size	256
Headend receiver FFT size	2048
Subchannel width	34.5 kHz
Symbol rate	32 kHz
Symbol period	31.25 μ s
Cyclic prefix length	2.26 μ s
Cyclic prefix length in samples	20
Maximum remote terminal sampling rate	8.832 MHz
Headend receiver sampling rate	70.656 MHz

segmented into eight disjoint *groups* of 128 consecutive subchannels. [10] Given this partitioning, any remote terminal can use a maximum of 128 subchannels that span 4.416 MHz bandwidth at any time. As a result, the maximum sampling rate of the remote terminals must be only 8.832 MHz, and the required FFT size is 256. Table 1 summarizes the example system parameters.

Each remote terminal configured according to the example parameters has a computational complexity that can be handled easily by a single DSP chip. Furthermore, even less expensive remote terminals can be designed. For some cost-sensitive applications, it may be desirable to design remote terminals that use fewer than 128 subchannels. For example, a video-conferencing signal may require 384 kbps. Assuming a minimum of 2 bits per subchannel, then at a symbol rate of 32 kHz only 6 subchannels are required to support this bit rate. To protect against subchannels that are too noisy or attenuated to support data, a conservative design for video-conferencing signals might use, say, 32 subchannels. The headend controller could then assign these smaller-bandwidth units to a *subgroup* within a particular group, as shown in Figure 6. The main advantages of using fewer than 128 subchannels are the resulting reduction in transmitter computational complexity and a reduced digital-to-analog converter complexity. When only 32 subchannels are used by a remote terminal, the required FFT size is 64, and the sampling rate, assuming the example configuration environment, is only 2.208 MHz. To maintain a subchannel spacing of 34.5 kHz, the cyclic prefix length is then 5 samples.

In addition to reducing remote terminal costs, the flexible eight-group configuration offers another advantage. As numerous cable operator studies have revealed, some parts of the reverse-channel bandwidth are poorer quality than other parts. In particular, it is generally agreed that the bandwidth from 5–20 MHz is significantly more susceptible to ingress than is the bandwidth from 20–42 MHz. [11] By designing the remote terminals with frequency-agile carriers, a remote terminal can transmit using the group of subchannels that is best able to support its data rate requirements. Hence, groups that can support high aggregate data rates can be used either by

some number of high bit rate users or by a larger number of lower bit rate remote terminals. Those groups that are affected by severe attenuation, passband ripple, or ingress can then be more lightly loaded. Furthermore, even though multiple independent DMT signals can be transmitted simultaneously in the eight-group configuration, a single headend receiver, spanning the entire reverse channel bandwidth, can be designed. The receiver performs a 2048-point FFT to demodulate incoming signals, and a controller then ensures signals are decoded correctly. If the complexity of the wideband receiver is too high for a particular installation, separate headend receivers can be designed for each of the eight groups.

5 Summary

In this paper, Synchronized Discrete Multi-Tone (SDMT) has been presented as a technique for coordinating multipoint-to-point reverse channel digital communications on cable or HFC networks. Although DMT has been applied traditionally only in point-to-point environments, this paper illustrates that, by appropriately synchronizing remote terminal transmissions, DMT can also be used effectively in multipoint-to-point environments. An example configuration was presented to show that a low-cost, low-complexity SDMT design can be implemented for cable or HFC networks. Thus, SDMT enables cost-effective multipoint-to-point reverse channel communications while preserving the well-known advantages of DMT, including its efficient use of the available bandwidth and its robustness to channel impairments and noise.

References

- [1] J.A.C. Bingham. "Multicarrier Modulation for Data Transmission : An Idea whose Time Has Come". *IEEE Communications Magazine*, 28(5):5–14, May 1990.




Figure 6: Partitioning of reverse channel into groups and subgroups

- [2] J.M. Cioffi. "A Multicarrier Primer". In *ANSI T1E1.4 Committee Contribution 91-157*, Boca Raton, FL, November 1991.
- [3] B.R. Saltzberg. "Performance of an Efficient Parallel Data Transmission System". *IEEE Transactions on Communication Technology*, COM-15:805-811, December 1967.
- [4] J. T. Aslanis. *Coding for Communication Channels with Memory*. PhD thesis, Stanford University, November 1989.
- [5] S. Kasturia. *Vector Coding for Digital Communication on Spectrally Shaped Channels*. PhD thesis, Stanford University, March 1989.
- [6] J.C. Tu. *Theory, Design and Application of Multichannel Modulation for Digital Communications*. PhD thesis, Stanford University, June 1991.
- [7] K.S. Jacobsen. *Discrete Multi-Tone-Based Communications on the Reverse Channel of Hybrid Fiber-Coax Networks*. PhD thesis, Stanford University, August 1996.
- [8] J.T. Aslanis, P.T. Tong, and T.N. Zogakis. "An ADSL Proposal for Selectable Forward Error Correction with Convolutional Interleaving". In *ANSI T1E1.4 Committee Contribution number 92-180*, August 1992.
- [9] J.M. Cioffi. "Discrete Multitone Data Transmission System Using an Overhead Bus for Synchronizing Multiple Remote Units". United States patent number 5,625,651, Issued April 29, 1997.
- [10] J.A.C. Bingham and K. Jacobsen. "CATV Reverse Channel Transmission Using Synchronized DMT,
- Part 1. Overview: Network, Data Rates and Services, Protocol". In *IEEE 802.14 Working Group Contribution number 95-001*, Boston, MA, January 1995.
- [11] Rogers Cable Systems. "Two-way Plant Characterization". 1994.