Multi-Carrier Modulation for Broadband Modems

Charles Sestok

December 12, 2000

1 Summary

Multi-carrier modulation (MCM) techniques for data transmission have been used in certain instances since the late 1950s [1]. In the 1990s, the international standard for Asynchronous Digital Subscriber Line (ADSL) broadband modems adopted Discrete Multitone (DMT) modulation, a multicarrier technique. Additionally, several wireless local area network standards have settled on Orthogonal Frequency Division Multiplexing (OFDM) as their modulation technique.

This summary focuses on DMT modulation for ADSL modems. After analyzing the asymptotic properties of multi-carrier systems in section 2, we'll examine a DMT modem complying with the ADSL standard. These modems implement channel equalization, mapping of bits to symbols, and channel coding using different techniques than traditional single carrier modems. Sections 3 and 4 will discuss the modulation and coding techniques used for ADSL modems. Finally, section 5 considers two weaknesses in DMT modems, susceptibility to radio-frequency (RF) interference, and a large peak-to-average power ratio.

2 Multi-Carrier Modulation Techniques

2.1 Idealized Multi-carrier Modulation Scheme

An MCM modem uses a set of parallel passband channels covering disjoint frequency bands to communicate data. The transmitter sums the $\bar{N}$ independent waveforms and the reciever separates them with frequency selective filtering. In a simplified situation, the subchannels use ideal bandpass filters with equal bandwidth, and equally spaced center frequencies to generate their waveforms. The bandwidth for each subchannel is Δf and the center frequencies for each channel are $f_b = b\Delta f$ for $b = 1, 2, \ldots, \bar{N}$. Each of these subchannels produces a QAM signal independent of the other subchannels. In the frequency domain, the transmitted signal for subchannel b is $X_b(f)$, and the recieved signal is $Y_b(f)$.

Asypmtotically, this idealized MCM system can achieve the capacity of a bandlimited, LTI communication channel without using filters to shape the transmit power spectrum or equalize the channel's impulse response. Consider an LTI channel with frequency response $H(f)$, and noise modeled by an additive white Gaussian noise process with power spectrum $U(f)$. If the total bandwidth used by the MCM modem is constant, $\Delta f \to 0$ as $\bar{N} \to \infty$. As the subchannel bandwidth approaches zero, the effect of $H(f)$ on each subchannel can be approximated by a constant gain, $Y_b(f) \approx H(f_b)X_b(f) + U(f)$. Asymptotically, each subchannel behaves like an ideal AWGN communication channel, and no equalization is required to optimally demodulate and detect the data. Additionally, the idealized MCM system can achieve the capacity-achieving transmit PSD without any spectral shaping by transmit filters. For a general channel, the "water-pouring" PSD achieves capacity. The MCM system can approximate the optimum PSD, or any other PSD, by varying the average energy assigned to each of the subchannels.

2.2 Performance Analysis of Multicarrier Modems

In a single carrier modem with QAM modulation and a fixed coding scheme, the channel's signal-to-noise ratio determines the achievable probability of error per dimension. The paper by Cioffi [2] reviews the performance analysis of single-carrier QAM and extends the formulation to MCM systems. It shows that an equivalent SNR, approximately equal to the geometric mean of the subchannel SNRs, characterizes the performance of an MCM modem.

For single-carrier QAM on ideal AWGN channels, the union bound estimate for the probability of a symbol error is

$$P_e \leq 4Q\left(\frac{d_{min}}{2\sigma}\right). \tag{1}$$

For ISI-free channels, the constellation minimum distance should be adjusted by the channel gain, $|H|$. This equation defines the minimum distance necessary for an uncoded system to communicate with a particular probability of error. An uncoded modem requires an SNR gap of 9.8 dB to operate at a symbol error probability of 10^{-7}. In many cases, a coding scheme is used to reduce the minimum distance necessary to achieve a desired probability of error. Additionally, the system may include an error margin, a safety factor included to protect the modem's performance in case of unanticipated channel degradation.

The SNR gap (alternatively referred to as normalized SNR) is a performance measure that accounts for both the coding gain and error margin in the performance analysis. The gap, Γ, is defined by

$$3\Gamma = \frac{d_{min}^2}{4\sigma^2}. \tag{2}$$

The union bound for the probability of error is $P_e \leq 4Q\left(\sqrt{3\Gamma}\right)$ in terms of the SNR gap. This formula can be used to analyze the impact of an error margin γ_m and a code with coding gain γ_c If an uncoded system requires an SNR gap of Γ_u to achieve the target error probability, the gap necessary with coding and error margin included is

$$\Gamma(dB) = \Gamma_u + \gamma_m - \gamma_c. \tag{3}$$

Additionally, the gap ties the channel SNR to the number of bits per symbol under the coding scheme. If the constellation has average energy E, the channel SNR is

$$SNR = \frac{E|H|^2}{2\sigma^2}.$$

In terms of the SNR and gap, the bits per symbol is

$$b = \log_2\left(1 + \frac{SNR}{\Gamma}\right). \tag{4}$$

In an MCM modem, the probability of error is the mean of the subchannel error probabilities. Thus, the largest subchannel error probability dominates the overall error probability expression. Typically, MCM modems are designed so that the subchannel error probabilities are equal. Since the idealized MCM modem can be treated as a set of parallel, ISI-free, AWGN channels, the single-carrier QAM error analysis can be applied. For an MCM system with $\bar{N}$ subchannels, each consisting of QAM signals with distances d_i, average energies E_i, channel gain H_i, and noise variance σ_i, the parameters determining the subchannel performance are

$$3\Gamma = \frac{|H_i|^2 d_i^2}{4\sigma_i^2} \tag{5}$$

$$SNR_i = \frac{|H_i|^2 E_i}{2\sigma_i^2} \tag{6}$$

$$b_i = \log_2\left(1 + \frac{SNR_i}{\Gamma}\right). \tag{7}$$

The total number of bits transported in one multi-carrier symbol is

$$b = \sum_{i=1}^{\bar{N}} \log_2 \left(1 + \frac{SNR_i}{\Gamma}\right).$$

The equation for b can be rewritten in terms of an "average" SNR in order to resemble the equation for single carrier QAM. In terms of the new SNR, the number of bits per MCM symbol is

$$b = \bar{N} \log_2 \left(1 + \frac{S\bar{N}R}{\Gamma}\right). \tag{8}$$

In terms of Γ and the subchannel SNRs, the average SNR is

$$S\bar{N}R = \Gamma \left(\left[\Pi_{i=1}^{\bar{N}} \left(1 + \frac{S\bar{N}R_i}{\Gamma}\right)\right]^{\frac{1}{\bar{N}}} - 1 \right).$$

In most applications, the average SNR can be approximated as the geometric mean of the subchannel SNRs. This performance figure can be used to characterize the performance of an MCM system and compare it directly to a single carrier modem.

3 Discrete Multitone Modulation - a practical modulation scheme

3.1 Discrete Multitone Modulation

Discrete Multitone (DMT) modulation is a discrete time approximation to the MCM system discussed in section 2.1. DMT uses the Discrete Fourier Transform (DFT) and Inverse DFT to demodulate and modulate the data to be communicated. The DFT and IDFT can be efficiently implemented with commercial DSPs using the Fast Fourier Transform (FFT). The increasing speeds of commercially available DSPs throughout the 1980s and 1990s made the implementation of DMT-based broadband modems practical. The international standard for ADSL adopted DMT in 1993 [3].

In DMT, the subcarriers, sometimes referred to as a "tones," are modulated into a real, discrete-time signal by an IFFT operation. The modulation is performed on a block-by-block basis. In block k, each tone produces a complex-valued QAM symbol $X_{i,k}$, where i ranges from 1 to $\bar{N} - 1$. A complex conjugate symmetry condition on the IFFT input ensures that the IFFT produces a real time-domain output block.The DC component is zero to maintain conjugate symmetry. Thus, a $2\bar{N}$ point IFFT produces a $2\bar{N}$ block of real time-domain samples, written as $x_{i,k}$ where i ranges from 0 to $2\bar{N} - 1$. This block is coverted into a waveform by a parallel-to-serial conversion and a D/A converter. If blocks arrive from the subchannels at a rate of $\frac{1}{T}$, the sampling rate of the DAC is $\frac{2\bar{N}}{T}$.

The data is demodulated in a similar procedure. The waveform is sampled by a A/D converter with rate $\frac{2\bar{N}}{T}$ and is collected into a block by serial-to-parallel conversion. An FFT demodulates the data and recovers the QAM symbols for each tone. The time-domain input to the demodulator is $y_{i,k}$, and the frequency-domain output is $Y_{i,k}$.

The IFFT and FFT are orthonormal transforms over $\mathbb{C}^{2\bar{N}}$. Each of the subcarriers can be treated independently of the others since they are modulated onto orthogonal signals. Additionally, demodulation by the FFT acts as a set of matched filters for the subcarriers.

3.2 Discrete Multitone over LTI Channels

When used over an LTI channel, however, the orthogonality property of DMT is lost. The channel frequency response distorts the power spectrum of the modulating vectors, causing interference between the subchannels. Fortunately, the orthogonality of the modulation vectors can be regained by introducing redundant samples into our data block. A cyclic prefix, a guard band of samples from the end of the data block can mitigate the effect of the channel's impulse response.

For example, consider an FIR channel with impulse response h_n, where $n = 0, \ldots, \nu$. The channel has frequency response $H(f)$, and its impulse response has a $2\bar{N}$ point DFT H_i. The circular convolution of two vectors in the time domain is equivalent to the product of their DFTs in the frequency domain. Thus, the circular convolution of h_n with $x_{i,k}$ would preserve the orthogonality of the DMT modulation vectors. In the absence of noise the recieved data would be

$$Y_{i,k} = H_i X_{i,k}$$

.

A cyclic prefix (CP) adds ν samples from the end of every block of data as a prefix. In this case, the kth data block becomes $x_{i,k}$, where $i = -\nu, \ldots, 0, 1, \ldots, 2\bar{N} - 1$. The samples with negative indices satisfy $x_{-i,k} = x_{2\bar{N}-i,k}$ for $i = 1, \ldots, \nu$. With the CP in place, the channel output also has $2\bar{N} + \nu$ samples per block. In the time domain, the channel output is $y_{i,k}, i = -\nu, \ldots, 2\bar{N} - 1$. If we drop samples with a negative index, the remaining data is the circular convolution of $x_{i,k}, i = 0, \ldots, 2\bar{N} - 1$ and $h_n, n = 0, \ldots, \nu$. The CP requires an increase in the system's bandwidth, but it restores the orthogonality of the modulation vectors for any channel with an impulse response less than $\nu + 1$ samples long.

3.3 DMT for Broadband Modems

The international standard for Asymmetric Digital Subscriber Line (ADSL) has adopted DMT as its modulation scheme. ADSL is a standard that allows for broadband access the home over existing copper wire in the telephone network. In the downstream direction, the bandwidth for transmitting data from the network to the home is 1.104 MHZ. In the upstream direction, the bandwidth is 138 KHz. The standard uses 255 subchannels for data transmission. The baseband subchannel is not used for data. The carrier separation is $\Delta f = 4.3125$KHz, and the data block rate is 4 KHz. The standard permits between 2 and 15 bits per subchannel symbol $X_{i,k}$ and requires 6 dB of error margin. The achievable data rates for DMT-based modems depend on the length of the telephone loop, and the modems can support rates greater than 8 Mb/s over wires less than 1 mile long.

4 Bit Loading and Coding for DMT Modems

In the ADSL application, a modem must operate over a wide variety of telephone channels. Fortunately, DMT can support highly adaptive bit allocation and coding schemes. The modem performs channel identification and uses estimates of the subchannel SNRs to assign bits to each tone. As argued earlier, the optimum allocation follows a "water-pouring" power spectrum. At a fixed Γ, the modem ought to send more data over tones with larger subchannel SNRs. In general, however, the water-pouring bit allocation does not produce bit allocations that are easily implemented with QAM constellations and coding schemes of reasonable complexity. The optimal bit allocations may require very fine precision in the number of bits per subchannel. Unfortunately, standard compliant ADSL modems are restricted to sending from 2 to 15 bits per symbol on each tone that transmits data. Practical coding schemes such as concatenated codes or turbo codes cannont produce infinitely fine granularity in bit allocations when combined with these constellations. A variety of bit loading algorithms have been proposed to generate good bit allocations with constellations of finite granularity [4], [5].

The ADSL standard uses sophisticated coding techniques to produce the required 10^{-7} error probability. Flexible coding schemes that adapt to the wide variety of potential bit allocations have been developed for use with DMT. Current versions of the standard exploit concatenated trellis and block codes [6], and the use of turbo trellis coding has been proposed for next-generation modems [7]. Using concatenated codes and trellis shaping, coding gains of 6.0 dB have been observed at $P_e \approx 10^{-7}$.

5 Impairments of DMT Modulation

DMT has become a viable modulation scheme for broadband data communication. It can be used to implement high performance modems that adapt to channel conditions and exploit sophisticated channel coding techniques. Though they has many desirable characteristics, DMT modems suffer from several important impairments that can degrade their performance. Two significant weaknesses of DMT modems are their susceptibility to RF interference, and the large peak-to-average power variation in DMT modulated signals.

5.1 RF Interference

The frequency band used by the ADSL standard extends to 1.104 MHz, and it overlaps the AM radio band. Some telephone channels will experience significant interference from the narrowband signals from nearby radio stations. Even though the copper wires in telephone lines are twisted, they can still act as an antenna if the RF interference is strong enough.

In the presence of a strong AM radio signal, the optimal PSD will not allocate data to frequencies occupied by the interference. The bit loading algorithms in section 3 can approximate this behavior for the idealized MCM system discussed in section 2.1. The modulation vectors are ideal bandpass filters, so only subcarriers overlapping the interference are affected. Unfortunately, DMT does not have perfect frequency selection. The FFT in the receiver can be viewed as a filterbank. It is well known that the frequency response of an FFT-derived filterbank has high sidelobes. Consider the subchannel with center frequency f_i. The frequency response magnitude of the subchannel's matched filter is only -13dB from the peak at $f_i \pm \Delta f$. The matched filter can pick up narrowband interference at adjacent tones, reducing the subchannel SNRs over a frequency band broader than the bandwidth of the interference. This effect will reduce the bit allocation over a wide band of tones, and does not lead to the optimal water-pouring PSD. A variety of techniques to mitigate narrowband interference have been examined, including windowing data at the reciever or modification of the modulation vectors to increase frequency selectivity [8].

5.2 Peak-to-Average Power Variation

A second major disadvantage of DMT modulation is the large amplitude variation of the transmitted signal. Through the central limit theorem, we can argue that the amplitude of the modulated waveform is well-modeled by a Gaussian random variable. The ratio of the peak to average signal amplitude, the peak-to-average ratio or PAR, is rather large. The signal may encounter nonlinear distortion such as clipping if it passes through amplifiers in analog portion of the modem that have limited dynamic range. Clipping of the signal acts as impulsive noise, correlated to the signal amplitude, and it affects all of the tones equally since the impulse has a flat PSD. A variety of approaches to reduce the PAR of DMT signals have been explored. Some PAR mitigation techniques use properties of the channel code to reduce the peak amplitude [9]. Alternative approaches have attempted to reduce PAR by adding time domain signals generated with unused tones to the DMT signal [10], [11]. These "dummy signals" carry no information, and are designed to reduce the peaks in the transmitted signal.

References

[1] J. A. C. Bingham. Multicarrier Modulation for Data Transmission: An Idea Whose Time Has Come. *IEEE Communications Magazine*, pp.5-14,May 1990.

[2] J. Cioffi. A Multicarrier Primer. In *ANSI T1E1.4 Committee Contribution 91-157.* Boca Raton, November 1991.

[3] T. Starr, J. Cioffi, and P. Silverman, *Understanding Digital Subscriber LineTechnology.* Englewood Cliffs, NJ:Prentice Hall, 1999.

[4] P. Chow, J. Cioffi, and J. Bingham. A Practical Discrete Multitone Transciever Loading Algorithm for Data Transmission over Spectrally Shaped Channels. *IEEE Trans. Comm.*, Vol. 43, No. 2/3/4, pp. 773-775, Feb./Mar./Apr. 1995.

[5] A. Leke and J. Cioffi. A Maximum Rate Loading Algorithm for Discrete Multitone Modulation Systems. *Proc. IEEE Globecomm 1997*, pp. 1514-1518.

[6] T. Z. Zogakis, J. Aslanis, and J. Cioffi. A Coded and Shaped Discrete Multitone System. *IEEE Trans. Comm*, Vol. 43, No. 12, pp. 2941-2949, Dec. 1995.

[7] J. Lauer and J. Cioffi. A Turbo Trellis Coded Discrete Multitone Transmission System. *Proc. IEEE Globecomm 1999*, Vol. 5, pp. 2581-2585.

[8] J. Bingham. RFI Suppression in Multicarrier Transmission Systems. *Proc. IEEE Globecomm 1996*, Vol. 2, pp. 1026-1030.

[9] V. Tarokh and H. Jafarkhani. On the Computation and Reduction of the Peak-to-Average Power Ratio in Multicarrier Communications. *IEEE Trans. Comm*, Vol. 48, No. 1, pp. 37-44, Jan. 2000.

[10] J. Tellado and J. Cioffi. Efficient Algorithms for Reducing PAR in Multicarrier Systems. *Proc. Intl. Symposium Info. Theory 1998*, pp. 191.

[11] A. Gatherer and M. Polley. Controlling Clipping Probability in DMT Transmission. *Proc. Asilomar Conf. on Computers and Communications, 1998*, Vol. 1, pp. 578-584.