

A Coded and Shaped Discrete Multitone System

T. Nicholas Zogakis, *Member, IEEE*, James T. Aslanis Jr., *Member, IEEE*,
and John M. Cioffi, *Senior Member, IEEE*

Abstract—In this paper, we show how coding and constellation shaping may provide significant gains to a discrete multitone (DMT) system transmitting over spectrally-shaped channels. First, we present and analyze a concatenated coding scheme consisting of an inner trellis code and outer block code when applied to DMT modulation, and we address some of the implementation issues associated with this scheme. Some laboratory test results for a DMT prototype employing the coding scheme are presented. Next, we propose a method for applying Forney's trellis shaper across the tones in a DMT system to realize significant shaping gain. To illustrate the coding and shaping gains achieved, we use scenarios indicative of the newly introduced asymmetric digital subscriber line service. By combining a powerful coding scheme, shaping, and DMT modulation, we arrive at an implementable transceiver that can provide very high data rates over spectrally-shaped channels.

I. INTRODUCTION

THE transmission of high-speed data over spectrally-shaped channels can be achieved by a sophisticated combined modulation and equalization technique. One approach that has been found well suited for a number of new applications is the use of multicarrier modulation, and we focus upon a particular form of multicarrier modulation known as discrete multitone (DMT) modulation [1]. With this approach, the channel is divided into a number of independent subchannels in the frequency domain, and bits and power are allocated among the subchannels to maximize throughput. The primary advantage of the multicarrier approach is that the problem of bandwidth optimization is greatly simplified, thus allowing very high data rates to be achieved with reasonable implementation complexity.

Given a DMT transceiver with the capability of performing bandwidth optimization, we address the problem of the application of coding and shaping to the system. First, we present a flexible, implementable, bandwidth-efficient coding scheme consisting of an outer block code and inner trellis code operating across the subchannels, and we analyze the performance of this scheme for a number of realistic scenarios. Implementation issues specific to DMT modulation

are addressed, and some laboratory test results are presented. Next, we show how trellis shaping may be applied across the tones in a DMT system to obtain significant shaping gain. By combining the proposed coding and shaping schemes, *real* (as opposed to asymptotic) coding plus shaping gains exceeding 6.0 dB at a bit error rate (BER) of 10^{-7} and 7.0 dB at a BER of 10^{-9} are achieved in the DMT system. To illustrate various results, we use examples from the asymmetric digital subscriber line (ADSL), a service proposed for providing a downstream data rate ranging from 1.544 Mbps to 6.4+ Mbps from the telephone company's central office to the customer, along with a lower-speed return channel over existing copper twisted pair [2].

In Section II, we review the concept of DMT modulation and establish the baseline DMT system to which we will apply both coding and shaping. In Section III, we show how coding is applied to the system and analyze the coding gains possible with the proposed scheme. Next, in Section IV, we address the problem of the application of shaping to DMT modulation, and we evaluate the shaping gains realized for the same scenarios used in the coding gain analysis of Section III. Finally, we summarize our results in Section V and conclude with the presentation of a practical multicarrier encoder structure.

II. BASELINE DMT SYSTEM

A simplified block diagram of the basic DMT transmitter is presented in Fig. 1 [3]. At the input to the system, the bit stream is partitioned into blocks of size $b = RT$ bits, where R is the uncoded bit rate, T is the DMT symbol period, and b is the number of bits contained in one DMT symbol. The bits collected during the m th symbol interval are allocated among $\bar{N}$ subchannels or tones in a manner determined during system initialization, with more bits given to those subchannels with higher signal-to-noise ratios (SNR's). Depending upon the data rate and SNR function, some of the channels may not be assigned any bits. We let b_i denote the number of bits assigned to the i th tone, so that $b = \sum b_i$. On subchannel i , the b_i bits, represented by the constellation label $d_{i,m}$, are mapped to a complex constellation point $X_{i,m} = f[d_{i,m}]$ in a constellation of size 2^{b_i} , where $f[\cdot]$ denotes the mapping operation. Next, a block of real time-domain samples, $\{x_{i,m}\}$, is formed by performing a length $N = 2\bar{N}$ Inverse Fast Fourier Transform (IFFT) on the complex symbols $\{X_{i,m}, i = 0, 1, \dots, \bar{N} - 1\}$, where $X_{0,m} = 0$ and $\{X_{i,m} = X_{N-i,m}^*, i = \bar{N} + 1, \bar{N} + 2, \dots, N - 1\}$. A cyclic prefix consisting of the last ν samples of the data block $\{x_{i,m}\}$ is added to the beginning of the block to form the signal transmitted over the channel. The receiver

Paper approved by E. Eleftheriou, the Editor for Equalization and Coding of the IEEE Communications Society. Manuscript received August 9, 1994; revised February 27, 1995.

T. N. Zogakis was with the Information Systems Laboratory, Stanford University, Stanford, CA 94305 USA. He is now with Amati Communications Corporation, Mountain View, CA 94040 USA (e-mail: zogakis@isl.stanford.edu).

J. T. Aslanis Jr. is with Amati Communications Corporation, Mountain View, CA 94040 USA (e-mail: redlands@isl.stanford.edu).

J. M. Cioffi is with the Information Systems Laboratory, Stanford University, Stanford, CA, 94305 USA (e-mail: cioffi@shannon.stanford.edu).

IEEE Log Number 9415929.

0090-6778/95\$04.00 © 1995 IEEE




Fig. 1. Uncoded/unshaped DMT transmitter.

corresponding to the transmitter illustrated in Fig. 1 merely consists of the inverse operations.

In this paper, we assume that $\bar{N}$ and ν are chosen sufficiently large so that the tones are well approximated as independent and memoryless subchannels.¹ Under these conditions, the complex point $Y_{i,m}$ obtained at the output of the i th subchannel is given by $Y_{i,m} = H_i X_{i,m} + N_{i,m}$, where H_i represents the complex gain of the subchannel and $N_{i,m}$ is a sample of a complex Gaussian noise process. Hence, the key to analyzing the uncoded and unshaped DMT system is to note that each of the $\bar{N}$ subchannels may be considered as supporting a low-speed, memoryless, quadrature amplitude modulated (QAM) signal [4]. We let $P_i = E\{|X_{i,m}|^2\}$ denote the two-dimensional (2-D) symbol power allocated to the i th subchannel and $2\sigma_i^2 = E\{|N_{i,m}|^2\}$ the 2-D noise variance. With these definitions, the output SNR for the i th tone is given by $\text{snr}_i = \frac{P_i |H_i|^2}{2\sigma_i^2}$.

A useful parameter for analyzing the DMT system's performance is the SNR gap, Γ , which represents the distance between the performance of a QAM signaling scheme and capacity over a memoryless channel [7], [8]. Using the SNR gap approximation, we can relate b_i to snr_i by the expression $b_i = \log_2(1 + \frac{\text{snr}_i}{\Gamma(C, P_{2-D})})$, where the SNR gap, $\Gamma(C, P_{2-D})$, is a function of the coding scheme C and 2-D error rate P_{2-D} . In formulating this relationship, we are assuming that the DMT system is designed to provide equal error rate on each subchannel.

For an uncoded DMT system operating at a 2-D error rate of P_{2-D} , the SNR gap may be written as $\Gamma_{\text{dB}}(C, P_{2-D}) = \Gamma_{0,\text{dB}}(P_{2-D}) + \gamma_{m,\text{dB}}$, where $\Gamma_{0,\text{dB}}(P_{2-D})$ is the nominal gap (in dB) required to meet the error rate and $\gamma_{m,\text{dB}}$ is any required system margin [4]. Using well-known QAM relationships, we can show that $P_{2-D} = 4Q(\sqrt{3\Gamma_0(P_{2-D})})$, where $Q(\cdot)$ is the Gaussian probability of error function. When a coding scheme with gain $\gamma_{c,\text{dB}}(C, P_{2-D})$ and shaping scheme with gain $\gamma_{s,\text{dB}}$ are applied, the SNR gap is reduced to $\Gamma_{\text{dB}}(C, P_{2-D}) = \Gamma_{0,\text{dB}}(P_{2-D}) + \gamma_{m,\text{dB}} - \gamma_{c,\text{dB}}(C, P_{2-D}) - \gamma_{s,\text{dB}}$ and thus $P_{2-D} = 4Q(\sqrt{3\Gamma(C, P_{2-D})\gamma_c(C, P_{2-D})\gamma_s/\gamma_m})$ [4], [8]. In the remainder of this paper, we present and analyze methods for applying coding and shaping to DMT modulation to achieve large $\gamma_{c,\text{dB}}(C, P_{2-D})$ and $\gamma_{s,\text{dB}}$ over spectrally-shaped channels.

¹ See [1], [4] for a discussion of the choice of $\bar{N}$ and [5], [6] for methods to reduce the length of ν .




Fig. 2. Coded DMT transmitter.




Fig. 3. Coded DMT receiver.

III. APPLICATION OF CODING

A. Coded DMT System Model

The coding scheme that we apply to the DMT system should have large coding gain, reasonable implementation complexity, flexibility (adaptable to data rate), and some measure of burst immunity. A suitable scheme for achieving all these goals is a concatenated coding scheme consisting of an inner trellis code operating across the tones as suggested in [9]–[11] and an outer block code with variable outer code parameters. The parameters of the outer code are determined according to the desired data rate and even may be determined on a channel by channel basis, if necessary. Figs. 2 and 3 illustrate how the proposed coding scheme is incorporated in the DMT system, where only those portions of the DMT transceiver involving the encoding of bits into complex symbols and the decoding of complex symbols into bits are depicted.

The binary data stream at the input to the system is first encoded by an outer, interleaved Reed–Solomon (RS) code with code length n and information length k . In traditional applications, block codes result in an increase in the symbol rate (i.e., increase in bandwidth) to maintain the same information rate, but for DMT modulation, the redundancy associated with the block code is absorbed by increasing the number of bits contained in each DMT symbol by a factor of n/k . As a result, the bit allocation algorithm makes the best tradeoff between bandwidth expansion, where a greater fraction of the $\bar{N}$ tones are used for transmission, and signal set expansion, where the sizes of the constellations supported by a subset of the used tones are increased, in distributing the additional bits.

In contrast to the block encoder, the inner trellis encoder in Fig. 2 operates on the bits at the output of the bit allocation unit, producing a set of complex symbols that serves as the input to the IFFT block. The only change from a basic trellis encoder as defined in [12] is that the signal selector component must select points from different size constellations as the encoder operates across the tones; the convolutional encoder and coset selector remain the same. As a result of the redundancy of the trellis code, the number of bits on the i th subchannel is increased to $b_i + \bar{r}$, where $\bar{r}$ is the normalized redundancy of the code in bits per 2-D symbol.

In the receiver, depicted in Fig. 3, a single Viterbi decoder operates across the subchannels on the noisy constellation points at the output of the FFT, which is not pictured. In using the Viterbi decoder across the subchannels, we are implicitly exploiting the important fact that the subchannels are independent and memoryless in the DMT system. Hence, the gain obtained from applying the trellis code will essentially be the same as that obtained in an intersymbol interference-free environment.

B. Theoretical Performance

To determine the theoretical coding gain afforded by the codes when applied to the DMT system, we consider the components of the overall gain individually. During the course of the analysis, it will be convenient to work with 2-D symbol error rates, bit error rates, and RS symbol error rates, depending upon what part of the system is being considered. For simplicity, we assume that these quantities are related by constant factors, and we use the 2-D error rate as a common basis. In particular, 2-D error rates are converted to bit error rates by multiplying by one-half. Similarly, 2-D error rates are converted to RS symbol error rates by multiplying by a constant c , where c represents the average number of tones contributing bits to each RS symbol.

With P_{bit} denoting the required bit error rate at the output of the overall system, the net coding gain is conveniently divided into the following components:

- 1) $\tilde{\gamma}_{\text{rs,dB}}((n, k), P_{\text{bit}})$, the gain obtained by the RS code *without* taking into account the penalty for the increased data rate.
- 2) $\gamma_{\text{tc,dB}}((n, k), P_{\text{bit}})$, the gain provided by the trellis code at the error rate required by the RS code to achieve P_{bit} .
- 3) $\gamma_{\text{loss,dB}}((n, k), b)$, the loss incurred for increasing the data rate.

The overall coding gain is computed as

$$\gamma_{\text{c,dB}}((n, k), P_{\text{bit}}) = \gamma_{\text{tc,dB}}((n, k), P_{\text{bit}}) + \tilde{\gamma}_{\text{rs,dB}}((n, k), P_{\text{bit}}) - \gamma_{\text{loss,dB}}((n, k), b). \quad (1)$$

Each of the coding gain components is now considered in turn.

$\tilde{\gamma}_{\text{rs,dB}}((n, k), P_{\text{bit}})$: A RS codeword is correctly decoded provided it contains no more than $t = \lfloor (n - k)/2 \rfloor$ errors [13]. Assuming sufficient interleaving so that the errors appear random at the RS decoder's input and assuming the RS decoder does not attempt to correct the codeword if greater than t errors are detected, we may relate the output RS symbol error rate P_{rs} to the input RS symbol error rate P_s by

$$\begin{aligned} P_{\text{rs}} &= \sum_{i=t+1}^n \frac{i}{n} \binom{n}{i} P_s^i (1 - P_s)^{n-i} \\ &= \sum_{i=t+1}^n \binom{n-1}{i-1} P_s^i (1 - P_s)^{n-i}. \end{aligned} \quad (2)$$

Hence, given the output bit error rate P_{bit} , we equate $P_{\text{rs}} = 2cP_{\text{bit}}$ and iteratively solve (2) for P_s . The corresponding input bit error rate is given by $P_b((n, k), P_{\text{bit}}) = P_s/(2c)$. This bit error rate is equated to the error rate at the output of a QAM

demodulator to obtain $P_b((n, k), P_{\text{bit}}) = 2Q(\sqrt{3\Gamma_{\text{rs}}})$, from which an expression for the SNR gap Γ_{rs} is derived.² Now, we define the SNR gap Γ_0 required for an uncoded system to achieve the output bit error rate P_{bit} as the solution to $P_{\text{bit}} = 2Q(\sqrt{3\Gamma_0})$. The gain of the RS code without taking into account the data rate penalty is computed as the difference

$$\tilde{\gamma}_{\text{rs,dB}}((n, k), P_{\text{bit}}) = \Gamma_{0,\text{dB}} - \Gamma_{\text{rs,dB}}. \quad (3)$$

$\gamma_{\text{tc,dB}}((n, k), P_{\text{bit}})$: In the concatenated coded system, the 2-D error rate required at the output of the trellis decoder is given by $P_{2-\text{D}} = 2P_b((n, k), P_{\text{bit}})$, where $P_b((n, k), P_{\text{bit}})$ is the bit error rate required at the input to the RS decoder to achieve P_{bit} . Given $P_{2-\text{D}}$, we determine the SNR gap $\Gamma_{\text{tc,dB}}(P_{2-\text{D}})$ required to achieve this error rate by using the trellis code's performance curve. The gain of the trellis code is given by

$$\gamma_{\text{tc,dB}}((n, k), P_{\text{bit}}) = \Gamma_{0,\text{dB}}(P_{2-\text{D}}) - \Gamma_{\text{tc,dB}}(P_{2-\text{D}}) \quad (4)$$

where $\Gamma_{0,\text{dB}}(P_{2-\text{D}})$ is the SNR gap for an uncoded system at the error rate $P_{2-\text{D}}$.

$\gamma_{\text{loss,dB}}((n, k), b)$: To determine the loss for the increased data rate associated with the RS code, we first define $P_{\text{tot}}^*(b)$ as the minimum amount of power required to achieve the data rate b in a DMT system with a gap of 0.0 dB. This quantity is determined as the solution to the following optimization problem

$$\begin{aligned} \min_{\{b_i \geq 0\}} P_{\text{tot}}(b) &= \sum_{i=1}^{\bar{N}} \frac{(2^{b_i} - 1)}{\text{snr}_{\text{ch},i}} \\ \text{subject to } b &= \sum_{i=1}^{\bar{N}} b_i, \end{aligned} \quad (5)$$

where the i th term in the first summation is understood to be equal to zero if $b_i = 0$ and $\text{snr}_{\text{ch},i} = 0$. In (5), $\{b_i\}$ represents the DMT bit distribution, and $\text{snr}_{\text{ch},i} = |H_i|^2/(2\sigma_i^2)$ is the channel gain-to-noise function. With $d(\cdot)$ denoting the permutation of the subchannel indices that results in a decreasing sequence for $\text{snr}_{\text{ch},d(i)}$, the solution to the optimization problem is given by

$$\begin{aligned} b_{d(i)} &= \begin{cases} \log_2(K \text{snr}_{\text{ch},d(i)}) & i \leq u \\ 0 & i > u \end{cases} \\ K &= 2^{\frac{1}{u} \left(b - \sum_{i=1}^u \log_2 \text{snr}_{\text{ch},d(i)} \right)}. \end{aligned} \quad (6)$$

The parameter u signifies the number of subchannels actually used for transmission and is determined by finding the largest integer $u \leq \bar{N}$ such that $\log_2(K \text{snr}_{\text{ch},d(u)}) > 0$. Hence, the power distribution is simply given by a discrete version of the well-known water-pour distribution

$$P_{d(i)} = \begin{cases} K - \frac{1}{\text{snr}_{\text{ch},d(i)}} & i \leq u \\ 0 & i > u \end{cases} \quad (7)$$

and the minimum power by

$$P_{\text{tot}}^*(b) = uK - \sum_{i=1}^u \frac{1}{\text{snr}_{\text{ch},d(i)}}. \quad (8)$$

²The factor of two in front of the $Q(\cdot)$ function results from our assumption that bit error rates are one-half the 2-D error rates.

The loss now may be expressed as

$$\gamma_{\text{loss,dB}}((n, k), b) = P_{\text{tot,dB}}^*(nb/k) - P_{\text{tot,dB}}^*(b). \quad (9)$$

C. Integer Constraints

The theoretical coding gain analysis provided in Section III.B assumes infinite bit granularity and allows infinite constellation sizes. In a practical DMT system, the bit granularity is determined by the complexity of the signal selector, while the constellation size is limited by finite precision, nonlinearity, and timing jitter constraints. Hence, we now consider the coding gains realized under more practical constraints on the bit distribution.

We assume that the number of bits assigned to each carrier must be an integer *before* the application of the trellis code. In addition, we restrict the number of bits per used tone to lie in the range $b_{\text{low}} \leq b_i \leq b_{\text{max}}$, where b_{low} and b_{max} are integers; typically, we have $b_{\text{low}} = 2$. The analysis remains the same as in Section III.B except for the computation of $\gamma_{\text{loss,dB}}((n, k), b)$, which is now given by

$$\gamma_{\text{loss,dB}}((n, k), b) = \hat{P}_{\text{tot,dB}}^*(nb/k) - \hat{P}_{\text{tot,dB}}^*(b) \quad (10)$$

where $\hat{P}_{\text{tot}}^*(b)$ is the solution to the following integer-constrained optimization problem

$$\begin{aligned} \min_{\{b_i\}} \hat{P}_{\text{tot}}^*(b) &= \sum_{i=1}^{\bar{N}} \frac{(2^{b_i} - 1)}{\text{SNR}_{\text{ch},i}} \\ \text{subject to } b &= \sum_{i=1}^{\bar{N}} b_i \\ b_i &\in \{0, b_{\text{low}}, b_{\text{low}} + 1, \dots, b_{\text{max}}\}. \end{aligned} \quad (11)$$

The solution to (11) is too difficult to compute in a practical DMT system. However, this optimization problem is closely related to the problem of determining the optimum bit and power allocation schemes for maximizing the margin of a DMT system at a given data rate. A practical, albeit ad-hoc, method for computing the DMT bit and power distributions is proposed in [14], and versions of this algorithm are implemented in current DMT products. It can be verified that for practical scenarios, the theoretical coding gains do not change as long as b_{max} is not too small. For instance, in the ADSL application $b_{\text{max}} = 10$ is usually sufficient, and $b_{\text{max}} = 15$ covers the worst scenarios.

D. Accommodation of Trellis Code Redundancy

Unlike the situation for traditional single-carrier systems, a trellis-coded DMT system that allows a fractional normalized redundancy, $\bar{r}$, must accommodate many different multidimensional constellation sizes. While these different constellations could be supported by using an algorithmic multidimensional encoder based on generalized cross constellations [15], [16] or some other multidimensional technique, the additional complexity may be unwarranted in a practical DMT system. Indeed, even with an integer bit assignment, a DMT system based on $2\bar{N}$ -point FFT's effectively supports a bit granularity of $1/\bar{N}$ bits per 2-D symbol.

Under the constraints that the bit distribution must be integer and that $b_i \leq b_{\text{max}}$ for *both* the cases of a system with and without trellis coding, we provide the following method by which a practical bit and power allocation algorithm that solves (11) can be used to determine the bit and power allocations for the trellis-coded DMT system:

- 1) Solve (11) for nb/k bits and find u , the number of subchannels required.
- 2) Compute $u_{\text{tc}} = \frac{\lceil \bar{r}u \rceil}{\bar{r}}$, the number of tones to be used in the trellis-coded case.
- 3) Sort the channel gain-to-noise ratios in descending order: $\text{SNR}_{\text{ch},d(\cdot)}$.
- 4) Set $\text{SNR}_{\text{ch},d(i)} = 0$, $i > u_{\text{tc}}$.
- 5) Set $b_{\text{tc}} = nb/k + \bar{r}u_{\text{tc}}$.
- 6) Solve (11) for b_{tc} given $\text{SNR}_{\text{ch},d(i)}$.

The suboptimal approach for accommodating the trellis code redundancy should result in approximately the same SNR loss as would be expected for expanding the constellation by a factor of $2^{\bar{r}}$ on each tone. However, an exception arises when many of the tones are already constrained by the limitation of b_{max} . Under these conditions, the redundant bits are forced onto poor subchannels, leading to a degradation in the gain expected for the trellis code. We refer to this effect as bit-capping. If $\hat{P}_{\text{tot,dB}}^*(nb/k)$ denotes the required power for the system without trellis coding and $\hat{P}_{\text{tot,dB}}^*(b_{\text{tc}})$ the required power for the system with trellis coding, the bit-capping loss incurred is given by

$$\gamma_{\text{bitcap,dB}} = \hat{P}_{\text{tot,dB}}^*(b_{\text{tc}}) - \hat{P}_{\text{tot,dB}}^*(nb/k) - \bar{r}10\log_{10}(2.0). \quad (12)$$

E. DMT Coding Gain Results

By using the analysis discussed above, we computed the coding gains expected at BER's of 10^{-7} and 10^{-9} for two ADSL scenarios. The first scenario corresponds to a frequency-division multiplexed system operating at 1.6 Mbps over ANSI loops [17] in the presence of near-end crosstalk (NEXT) from 49 digital subscriber line (DSL) disturbers. The second is for the case of an echo canceled system transmitting at 6.4 Mbps over CSA loops [18] in the presence of NEXT from 10 DSL and 24 high bit-rate DSL (HDSL) disturbers, and NEXT and far-end crosstalk (FEXT) from 10 ADSL disturbers. The configurations of the eight specific loops used in the analysis are given in Fig. 4, where the two numbers above each segment give the length of the segment in feet and the gauge of the wire (AWG). Based on the spectral efficiencies typically required for the two data rates along with simulation results, we set $c = 2.5$ for the 1.6 Mbps scenario and $c = 1.5$ for the 6.4 Mbps scenario. The parameters common to both scenarios include a RS information length of $k = 200$, the use of 512-length FFT's, a 4.0 kHz DMT symbol rate, bit assignment constraints of $b_{\text{low}} = 2$ and $b_{\text{max}} = 10$, a 4.3125 kHz carrier spacing, and the presence of additive white Gaussian noise (AWGN) with a two-sided power spectral density of -143.0 dBm/Hz.

Table I presents the maximum coding gains achieved at BER's of 10^{-7} and 10^{-9} for both the cases of an applied RS code over GF(256) and a concatenated code consisting of an




Fig. 4. Loops used for evaluation.

TABLE I
THEORETICAL CODING GAINS FOR BER'S OF 10^{-7} AND 10^{-9}

loop, rate (Mbps)	10^{-7} error rate				10^{-9} error rate			
	RS		RS+Wei		RS		RS+Wei	
	$\gamma_{c,dB}$	n	$\gamma_{c,dB}$	n	$\gamma_{c,dB}$	n	$\gamma_{c,dB}$	n
ANSI 1, 1.6	3.5	226	5.4	212	4.3	226	6.4	212
ANSI 6, 1.6	3.4	226	5.2	210	4.3	226	6.2	214
ANSI 9, 1.6	3.7	228	5.4	210	4.5	234	6.4	212
ANSI 15, 1.6	3.6	226	5.4	212	4.5	228	6.4	214
CSA 1, 6.4	3.1	216	5.2	208	3.9	220	6.2	208
CSA 2, 6.4	3.1	214	5.2	206	3.8	216	6.1	208
CSA 4, 6.4	3.1	216	5.2	208	3.9	218	6.1	210
CSA 6, 6.4	3.2	216	5.3	208	3.9	218	6.2	208

outer RS code and an inner 16-state, 4-D trellis code devised by Wei [15]. The performance of the Wei code was obtained by combining Monte Carlo simulation results for high error rates and distance spectrum analysis for low error rates to generate an accurate curve over a wide range of error rates [19]. In obtaining the DMT coding gains, we used integer bit distributions and accommodated the trellis code redundancy in the suboptimal fashion described in Section III.D. Also listed in Table I is the optimum code length for realizing each of the coding gains listed. From the results, we find that RS codes can provide over 3.0 dB of gain at a BER of 10^{-7} , while the concatenated code provides over 5.0 dB of gain at the same error rate. We note that the gains obtained for a RS(216,200) code, the default code specified in the ADSL standard for the two data rates considered, differ from the optimum gains listed in Table I by at most 0.35 dB.

To confirm our coding gain analysis, we conducted laboratory tests on a DMT prototype for ADSL in October, 1993 at Amati Communications Corporation. The Amati DMT system operates at a sampling rate of 2.208 MHz and a DMT symbol rate of 4.0 kHz. A length 512 FFT is used for modulation and demodulation in the downstream direction, resulting in a carrier spacing of 4.3125 kHz. Furthermore, a maximum of $b_{\max} = 11$ bits per tone is supported. We tested performance at an effective data rate of 6.208 Mbps over a 6 kft, 26 AWG loop in the presence of AWGN with two different RS codes, a RS(202,194) code and a RS(210,194) code. Plots of the data points obtained in the lab are presented in Fig. 5 along with solid error rate curves derived using the same analysis techniques as those used to obtain the results in Table I. The theoretical curve and set of data points on the far right of the graph correspond to the uncoded system, while the group of




Fig. 5. Laboratory results for case of AWGN and a 6 kft, 26 AWG loop.

curves on the far left correspond to the concatenated RS + Wei coded system. The middle group of curves pertain to the case where only a RS code was applied to the data stream.

The correspondence between the data points and the theoretical performance curves is quite good, thus verifying the accuracy of our analysis. We note that the apparent coding gains in Fig. 5 are larger than any of the gains in Table I because the "uncoded" curve in Fig. 5 is actually for an uncoded system at the *same data rate* as the RS(202,194) coded system. In other words, all the curves in the figure that correspond to system performance with coding do not include the penalty for the addition of eight check bytes.³ The penalty is easily determined to be 0.8 dB in this case, and by subtracting 0.8 dB from the coding gains implied by Fig. 5, we find that the extrapolated gains are 3.1 dB for the RS(210,194) code and 5.2 dB for the RS(202,194) + Wei code at a BER of 10^{-7} .

IV. APPLICATION OF SHAPING

A. Shaping Across the Tones

Now that a good coding scheme for achieving large $\gamma_{c,dB}(C, P_{2-D})$ in a DMT system has been derived and analyzed, we turn to the problem of the application of shaping to achieve large $\gamma_{s,dB}$. As in the case of trellis coding, a straightforward approach would be to shape the constellations on each of the subchannels individually, but this would require too much delay, storage, and complexity. Hence, we consider a method of applying trellis shaping [20] *across* the tones in the DMT system. For simplicity, we focus only on the 4-state shaper discussed in [20], and we concentrate on the case in which one large circular constellation is stored in memory with an embedded labeling scheme to support the various size constellations needed on the DMT subchannels. The embedded labeling scheme is generated by ordering points on the half-integer grid $Z^2 + (0.5, 0.5)$ according to increasing energy

³The curves involving a RS(210,194) code do include the penalty for increasing from 6.464 Mbps to 6.720 Mbps.




Fig. 6. The 4-state, 2-D trellis shaper applied across the tones.

and assigning larger labels to points with higher energy, where labels are interpreted as unsigned integers. For the interested reader, an algorithmic encoder that generates square and cross constellations with labeling schemes that work well with the 4-state shaper is provided in [21].

Fig. 6 illustrates the trellis shaper that is applied across the tones. The first step in applying the shaper is to double the size of the constellation on each tone. Next, the sequence of most-significant bits $\{d_{i,m}^{b_i-1}\}$ of the unshaped constellation labels is passed through a rate one-half coset representative generator $(H_s^{-1}(D))^T$ to form an initial sequence of region specifier bits $\{v_{i,m}^{b_i}, v_{i,m}^{b_i-1}\}$ for the expanded constellations. Since an embedded labeling scheme is used, the region specifier bits $(v_{i,m}^{b_i}, v_{i,m}^{b_i-1})$ produced during symbol period m designate one of four concentric circular rings. This concept is depicted in the left-hand side of Fig. 7, where the four regions are labeled with the four possible values of $(v_{i,m}^{b_i}, v_{i,m}^{b_i-1})$. The sequence $\{v_{i,m}^{b_i}, v_{i,m}^{b_i-1}\}$ is subsequently modified by the modulo two addition of a codeword $\{c_{i,m}^1, c_{i,m}^0\}$ from the rate one-half convolutional code C_s with generator matrix $G_s(D) = [1 + D^2 \ 1 + D + D^2]$ and parity check matrix $H_s(D)$ to form the sequence $\{z_{i,m}^{b_i}, z_{i,m}^{b_i-1}\} = \{v_{i,m}^{b_i} \oplus c_{i,m}^1, v_{i,m}^{b_i-1} \oplus c_{i,m}^0\}$. The choice of codeword is determined by performing the Viterbi algorithm to search for the lowest-energy sequence out of an equivalence class of possible transmit sequences [20]. In particular, the squared magnitude of the constellation point

$$\sqrt{\frac{P_i}{P_o(b_i)}} f_o[v_{i,m}^{b_i} \oplus c_{i,m}^1, v_{i,m}^{b_i-1} \oplus c_{i,m}^0, d_{i,m}^{b_i-2}, \dots, d_{i,m}^1, d_{i,m}^0] \quad (13)$$

may serve as the weight of a branch in the trellis for the shaping code, where the branch is associated with the label $(c_{i,m}^1, c_{i,m}^0)$. In (13), $f_o[\cdot]$ represents the mapping of labels to points in the stored constellation, and $P_o(b_i)$ is the average 2-D symbol power of a 2^{b_i} -point constellation based on the half-integer grid. Hence, division by $\sqrt{P_o(b_i)}$ normalizes the unshaped constellation to unity energy, while multiplication by $\sqrt{P_i}$ incorporates the DMT power distribution. Once a trellis path is chosen, a mapping operation that includes transmit power scaling is performed to obtain the transmitted points $\{X_{i,m}\}$.

Before assessing the performance of shaping across the tones in a DMT system, it is instructive first to evaluate the performance of the 4-state shaper for different fixed constellation sizes. Simulation results reported in [20] show that a gain on the order of 1.0 dB with respect to a "square" constellation is obtained for the case of a 128-point constel-




Fig. 7. Illustration of standard and modified labeling schemes.




Fig. 8. Performance of 4-state trellis shaping code as a function of constellation size.

lation. However, it is not clear whether such gains may be achieved for small constellations for which the continuous approximation [16] does not hold. To ascertain the relationship between constellation size and shaping gain, simulations were conducted using constellation sizes ranging from 4 to 1024 points, and the results are presented in Fig. 8 for the shaping gains achieved with respect to the original, unshaped, circular constellation. The lower curve in Fig. 8 corresponds to the case in which the original embedded labeling scheme was maintained, whereas the upper curve was achieved when the labeling scheme was modified as discussed in [20] such that the points in the outer-most and next-to-inner-most rings were labeled from high energy to low energy and the two most-significant bits of each constellation label in the two middle rings were interchanged. The latter modification is equivalent to interchanging our definitions of regions one and two. Fig. 7 compares the standard embedded labeling scheme (left-hand side) with the modified labeling scheme (right-hand side) for an arbitrary constellation size, where the arrows indicate the direction of increasing constellation labels within each region.

The modified labeling scheme is not compatible with the embedded labeling scheme that would be used to store a constellation in the DMT system. Moreover, creating different look-up tables for all possible constellation sizes is clearly impractical. Fortunately, the desired labeling scheme may be induced on a stored, embedded constellation by using the following algorithm to modify the signal selector.

Given an l -bit constellation label, $(d^{l-1}, d^{l-2}, d^{l-3}, \dots, d^2, d^1, d^0)$, first check to see whether or not $l \leq 4$. If so, then the constellation point X is obtained in the normal way: $X = f_o[d^{l-1}, d^{l-2}, d^{l-3}, \dots, d^2, d^1, d^0]$, where $f_o[\cdot]$ denotes the original mapping operation. On the other hand, if $l > 4$, then the mapping is performed based on the values of the region specifier bits (d^{l-1}, d^{l-2}) as follows

$$\begin{aligned} d^{l-1}d^{l-2} = 00 &\Rightarrow X = f_o[d^{l-1}, d^{l-2}, d^{l-3}, \dots, d^2, d^1, d^0], \\ d^{l-1}d^{l-2} = 01 &\Rightarrow X = f_o[\bar{d}^{l-1}, \bar{d}^{l-2}, d^{l-3}, \dots, d^2, d^1, d^0], \\ d^{l-1}d^{l-2} = 10 &\Rightarrow X = f_o[\bar{d}^{l-1}, \bar{d}^{l-2}, \bar{d}^{l-3}, \dots, \bar{d}^2, d^1, d^0], \\ d^{l-1}d^{l-2} = 11 &\Rightarrow X = f_o[d^{l-1}, d^{l-2}, \bar{d}^{l-3}, \dots, \bar{d}^2, d^1, d^0] \end{aligned}$$

where $\bar{d}$ is the binary complement of d . Inverse operations may be performed in the receiver to retrieve the original constellation labels.

The dependency of the achievable shaping gain on the size of the constituent constellation implies that the overall shaping gain will depend on the DMT scenario and in particular on the bit and power distributions. If a separate trellis shaper were applied to each tone, then we could calculate the shaping gain as

$$\gamma_{s,\text{tot}} = \frac{\sum_{i \in \mathcal{U}} P_i}{\sum_{i \in \mathcal{U}} \frac{P_i}{\gamma_s(b_i)}} \quad (14)$$

where $\gamma_s(b_i)$ represents the dependency of the shaping gain on the constellation size as illustrated in Fig. 8. The summations in (14) are taken over the set of indices $\mathcal{U}$ corresponding to tones actually used for transmission. However, the shaping gain achieved when applying shaping across the tones will not be exactly the same as for separate trellis shapers, and the gain must be obtained through simulation. In fact, we claim that the gain will always be better than for separate shapers, whenever the power distribution is nonflat. This assertion is difficult to prove for low numbers of bits per 2-D symbol, although it has been verified through simulations. A theoretical justification based on the continuous approximation, which is valid for large numbers of bits per 2-D symbol, is given in Appendix A.

B. Coding and Shaping

To this point, we have considered the application of shaping to DMT modulation without regard to our applied coding scheme. The block code and applied shaping scheme are nearly separable, but for the trellis code, the shaper must not violate the sequence of least-significant bits of the constellation labels. In the case of the 4-state trellis shaper, two possible solutions for ensuring compatibility with a mod-2 trellis code are presented in [21], where the assumption is made that the DMT system accommodates the redundancy of the trellis code by using the algorithm proposed in Section III.D to achieve an integer bit distribution. The first solution is to simply exclude carriers supporting only two bits from the shaping operation, while the second is to prune branches (i.e., limit allowable transitions) from the shaping trellis at transitions corresponding to 4-point carriers so that the least-significant bits are not violated on any possible path. The advantage of

TABLE II
PERFORMANCE OF TRELLIS SHAPING FOR DMT EXAMPLES

loop, rate (Mbps)	γ_{sAB} at 10^{-7}			γ_{sAB} at 10^{-9}		
	separate	down block	w.r.t. sq/cr	separate	down block	w.r.t. sq/cr
ANSI 1, 1.6	0.81	0.87	0.90	0.81	0.87	0.90
ANSI 6, 1.6	0.99	1.07	1.10	0.98	1.07	1.09
ANSI 9, 1.6	0.80	0.85	0.86	0.84	0.88	0.89
ANSI 15, 1.6	0.90	0.96	0.98	0.88	0.93	0.95
CSA 1, 6.4	0.89	0.92	0.99	0.89	0.92	0.99
CSA 2, 6.4	0.84	0.86	0.96	0.84	0.87	0.97
CSA 4, 6.4	0.81	0.84	0.92	0.82	0.85	0.92
CSA 6, 6.4	0.82	0.86	0.94	0.82	0.86	0.94

the second solution from an implementation standpoint is that the regularity of the Viterbi operation is not altered. Both approaches yield approximately the same shaping gains.

C. DMT Shaping Gain Results

To evaluate the approach to shaping for DMT modulation discussed in Section IV.A, we used the same scenarios as investigated in Section III.E. Shaping gains were obtained for the bit and power distributions corresponding to the optimal concatenated code parameters listed in Table I. The constellation expansion associated with the trellis shaper was accommodated by allowing a maximum constellation size of 2048 points after shaping. For simplicity, trellis code compatibility was maintained by excluding carriers supporting two bits from the shaping operation.

Table II presents the shaping gains obtained with the 4-state trellis shaper for the two error rates considered earlier. The gains in the columns labeled as *separate* were obtained by applying an independent shaper to each tone, and those listed in columns labeled as *down block* were achieved by performing shaping down the DMT symbol. The gains in these four columns indicate the reduction in transmit power afforded by the shaping technique when measured relative to the original, unshaped signaling scheme. Additional columns labeled *w.r.t. sq/cr* are included in Table II to illustrate the shaping gains achieved for the case of shaping across the tones when measured relative to a DMT system that uses 2-D squares and crosses on the subchannels. This normalization is useful for comparing DMT systems that employ different signal selectors. The results in Table II show that trellis shaping across the tones provides anywhere from 0.85–1.10 dB of shaping gain for these scenarios. In addition, the down-the-block approach is found to perform slightly better than when independent shapers are applied, as expected.

V. SUMMARY

Fig. 9 depicts the encoder structure at which we finally arrive by putting all the components considered in this paper together. Combining the results provided in Tables I and II to form Table III, we find that over 6.0 dB of coding plus shaping gain is achieved at a BER of 10^{-7} and 7.0 dB at 10^{-9} .

The encoder in Fig. 9 is both high-performing as well as implementable. In addition, it should be noted that the encoder is not specific to DMT modulation; the only function of the




Fig. 9. Coded and shaped multicarrier encoder.

TABLE III
COMBINED CONCATENATED CODING PLUS SHAPING GAINS FOR DMT EXAMPLES

loop, rate (Mbps)	$\gamma_{cAB} + \gamma_{sAB}$ at 10^{-7}	$\gamma_{cAB} + \gamma_{sAB}$ at 10^{-9}
ANSI 1, 1.6	6.3	7.3
ANSI 6, 1.6	6.3	7.3
ANSI 9, 1.6	6.3	7.3
ANSI 15, 1.6	6.4	7.3
CSA 1, 6.4	6.1	7.1
CSA 2, 6.4	6.1	7.0
CSA 4, 6.4	6.0	7.0
CSA 6, 6.4	6.2	7.1

structure is to convert sequences of bits into $\bar{N}$ sequences of complex symbols. For this reason, we refer to the encoder in Fig. 9 as a multicarrier encoder. Given this structure, all that we need is some method for converting a spectrally-shaped transmission channel into a number of parallel, independent, memoryless subchannels; and a variety of approaches for accomplishing this partitioning have been proposed in the literature [22], [23]. However, DMT modulation currently provides the most efficient means for achieving the channel partitioning. Hence, by combining the encoder of Fig. 9 with DMT modulation, we arrive at an implementable transceiver structure that comes closer to the capacity of spectrally-shaped channels in practice than any other structure of which we know.

APPENDIX A OPTIMAL SHAPING FOR NONFLAT POWER DISTRIBUTION

In this Appendix, we give a theoretical argument as to why trellis shaping across the tones performs better than trellis shaping on a tone by tone basis, whenever the power distribution is nonflat. Our assertion is based on the observation that selecting points from an optimal $2N$ -dimensional region in each 2-D coordinate is not optimal when compared to selecting points from an optimal $2N$ -dimensional region that takes into account the different scaling in each of the coordinates.⁴ We state this observation as the following theorem.

Theorem 1 [Optimal Shaping for a Nonflat Power Distribution:] An optimized $2N$ -dimensional shaper operating down the block will provide higher performance than when

⁴Note that we place no constraints on complexity or constellation expansion in defining optimal.

optimized on a tone by tone basis, whenever the power distribution is nonflat.

Proof: Consider the optimal selection of points from a $2N$ -dimensional region. Let $\{b_i, i = 1, \dots, N\}$ and $\{P_i, i = 1, \dots, N\}$ denote the initial bit and power allocation schemes over N 2-D coordinates based on square constellations. By the continuous approximation, $P_i = 2^{b_i} d_i^2 / 6$, where d_i is the minimum distance between constellation points on the i th subchannel. Hence, the average transmit power per 2-D symbol is given by

$$P_{\text{sqr}} = \frac{1}{N} \sum_{i=1}^N P_i = \frac{1}{N} \sum_{i=1}^N 2^{b_i} d_i^2 / 6. \quad (15)$$

When a $2N$ -dimensional shaper is applied to each tone separately, the optimum scheme is to choose points in each subchannel from a $2N$ -dimensional sphere, resulting in an average transmit power of

$$P_{\text{sep}} = P_{\text{sqr}} / \gamma_{s\otimes N} \quad (16)$$

where $\gamma_{s\otimes N} = (\pi(N+1)) / (6(N!)^{1/N})$ is the shaping gain of a $2N$ -sphere [16]. On the other hand, when shaping is applied down the block, the points are selected from a sphere on a lattice scaled by d_i^2 in each 2-D coordinate. Letting R denote the radius of the sphere, we have

$$2^\beta = \left(\frac{V_{\otimes N}(R)}{\Delta V} \right)^{1/N} = \frac{\pi R^2}{(N!)^{1/N} \left(\prod_{i=1}^N d_i^2 \right)^{1/N}} \quad (17)$$

where $\Delta V = \prod_{i=1}^N d_i^2$ is the fundamental volume, $\beta = \frac{1}{N} \sum_{i=1}^N b_i$ is the normalized data rate, and $V_{\otimes N}(R) = (\pi^N / N!) R^{2N}$ is the volume of a $2N$ -dimensional sphere of radius R . The average power of a $2N$ -sphere is $P_{\otimes N}(R) = R^2 / (N+1)$, and thus the average power required is

$$P_{\text{blk}} = \frac{2^\beta (N!)^{1/N} \left(\prod_{i=1}^N d_i^2 \right)^{1/N}}{\pi(N+1)} = \frac{1}{\gamma_{s\otimes N}} \left(\prod_{i=1}^N P_i \right)^{1/N} \quad (18)$$

Taking the ratio of (16) and (18), we find that the advantage of down-the-block is given by

$$\frac{P_{\text{sep}}}{P_{\text{blk}}} = \frac{\frac{1}{N} \sum_{i=1}^N P_i}{\left(\prod_{i=1}^N P_i \right)^{1/N}} \quad (19)$$

Since the arithmetic average of a set of non-negative numbers is always greater than or equal to the geometric average, $P_{\text{sep}} / P_{\text{blk}} \geq 1$ with equality obtained when $P_i = P_j \forall i, j$. $\square$

REFERENCES

- [1] J. A. C. Bingham, "Multicarrier modulation for data transmission: An idea whose time has come," *IEEE Commun. Mag.*, vol. 28, no. 5, pp. 5-14, May 1990.
- [2] W. Y. Chen and D. L. Waring, "Applicability of ADSL to support video dial tone in the copper loop," *IEEE Commun. Mag.*, vol. 32, no. 5, pp. 102-109, May 1994.
- [3] P. S. Chow, J. C. Tu, and J. M. Cioffi, "Performance evaluation of a multichannel transceiver system for ADSL and VHDLS services," *IEEE J. Select. Areas Commun.*, vol. 9, no. 6, pp. 909-919, Aug. 1991.
- [4] J. M. Cioffi, "A multicarrier primer," ANSI T1E1.4 Committee Contribution, no. 91-157, Boca Raton, FL, Nov. 1991.

- [5] J. S. Chow, J. M. Cioffi, and J. A. C. Bingham, "Equalizer training algorithms for multicarrier modulation systems," in *1993 Int. Conf. Communications*, 1993, pp. 761–765.
- [6] J. S. Chow, "Finite-length equalization for multi-carrier transmission systems," Ph.D. dissertation, Stanford University, June 1992.
- [7] J. M. Cioffi, G. P. Dudevoir, M. V. Eyuboğlu, and G. D. Forney Jr., "MMSE decision-feedback equalizers and coding—part II: Coding results," *IEEE Trans. Commun.*, vol. 43, no. 10, pp. 2595–2604, Oct. 1995.
- [8] G. D. Forney Jr. and M. V. Eyuboğlu, "Combined equalization and coding using precoding," *IEEE Commun. Mag.*, vol. 29, no. 12, pp. 25–34, Dec. 1991.
- [9] A. Ruiz and J. M. Cioffi, "A frequency-domain approach to combined spectral shaping and coding," in *1987 Int. Conf. Communications*, 1987, pp. 1711–1715.
- [10] J. C. Tu and J. M. Cioffi, "A loading algorithm for the concatenation of coset codes with multichannel modulation methods," in *1990 Global Telecommunications Conf.*, 1990, pp. 1183–1187.
- [11] J. C. Tu, "Theory, design, and application of multi-channel modulation for digital communications," Ph.D. dissertation, Stanford University, June 1991.
- [12] G. D. Forney Jr., "Coset codes I: Introduction and geometrical classification," *IEEE Trans. Inform. Theory*, vol. 34, no. 5, pp. 1123–1151, Sept. 1988.
- [13] R. E. Blahut, *Theory and Practice of Error Control Codes*. Menlo Park, CA: Addison-Wesley, 1983.
- [14] P. S. Chow, J. M. Cioffi, and J. A. C. Bingham, "A practical discrete multitone transceiver loading algorithm for data transmission over spectrally shaped channels," *IEEE Trans. Commun.*, vol. 43, no. 3, pp. 773–775, Mar. 1995.
- [15] L. F. Wei, "Trellis-coded modulation with multidimensional constellations," *IEEE Trans. Inform. Theory*, vol. 33, no. 4, pp. 483–501, July 1987.
- [16] G. D. Forney Jr. and L.-F. Wei, "Multidimensional constellations—part I: Introduction, figures of merit, and generalized cross constellations," *IEEE J. Select. Areas Commun.*, vol. 7, no. 6, pp. 877–892, Aug. 1989.
- [17] K. Sistanizadeh, "Proposed canonical loops for ADSL and their loss characteristics," ANSI T1E1.4 Committee Contribution, no. 91–116, Aug. 1991.
- [18] *Belcore Technical Advisory TA-NWT-001210*, Generic Requirements for High-Bit-Rate Digital Subscriber Lines. Issue 1, Nov. 1991.
- [19] T. N. Zogakis, J. T. Aslanis Jr., and J. M. Cioffi, "Analysis of a concatenated coding scheme for a discrete multitone modulation system," in *1994 IEEE Military Communications Conf.*, 1994, pp. 433–437.
- [20] G. D. Forney Jr., "Trellis shaping," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 281–300 Mar. 1992.
- [21] T. N. Zogakis, "Coding and shaping for discrete multitone modulation over band-limited channels," Ph.D. dissertation, Stanford University, June 1994.
- [22] S. Kasturia, J. Aslanis, and J. M. Cioffi, "Vector coding for partial-response channels," *IEEE Trans. Inform. Theory*, vol. 36, no. 4, pp. 741–762, July 1990.
- [23] M. A. Tzannes, M. C. Tzannes, and H. Resnikoff, "The DWMT: A multicarrier transceiver for ADSL using M-band wavelet transforms," ANSI T1E1.4 Committee Contribution, no. 93–67, Miami, FL, Mar. 1993.



T. Nicholas Zogakis (S'94–M'95) received the B.S. degree in electrical engineering from the University of Florida, Gainesville, FL, in 1989, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA in 1990 and 1994, respectively.

During the 1990–1991 academic year, he was employed at Harris Corporation in Palm Bay, FL, where he worked in the area of signal identification and modulation recognition. Since September 1994, he has been employed at Amati Communications Corporation in Mountain View, CA, where he works on the analysis, design, and implementation of DMT modems. His research interests include digital communications, modulation, and coding theory.

Dr. Zogakis is a member of Eta Kappa Nu, Phi Kappa Phi, and the IEEE Communications and Information Theory Societies.

James T. Aslanis Jr. (S'88–M'89) received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, in 1981, the M.S. degree in systems engineering from Case Western Reserve University, Cleveland, OH, in 1983, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1990.

From 1983 to 1985, he was a Member of Technical Staff in the Switching Systems Engineering Laboratory, Holmdel, NJ, at AT&T Bell Laboratories. From 1990 to 1992, he was an Assistant Professor of Electrical Engineering at the University of Hawaii at Manoa, Honolulu, HI. In 1992, he joined Amati Communications Corp., Mountain View, CA, where he manages the design and development of transmission products employing discrete multitone technology. His research interests include digital communications, coding and information theory.

John M. Cioffi (S'77–M'78–SM'91), for a photograph and biography, see p. 2594 of the October issue of this TRANSACTIONS.