Combined Equalization and Coding Using Precoding

Techniques for approaching the capacity of high-SNR bandlimited channels have been developed only in the past decade. Precoding is an effective equalization method for such channels, and is now in use in telephone-line modems.

by G. David Forney, Jr., and M. Vedat Eyuboğlu

Shannon's formula for the capacity of a bandlimited Gaussian channel was one of the earliest results in information theory [1], and remains the most famous. This remarkable theorem showed that in principle enormous improvements were possible in the performance of digital communications systems. It led immediately to intense research aimed at the development of practical schemes that could begin to approach this fundamental limit.

The initial era of excitement was succeeded by an era of discouragement, as such practical schemes proved difficult to find. This feeling was summarized in a well-known folk theorem [2] which has been paraphrased as: "All codes are good, except those that we know of."[1] Communications engineers tended to regard the channel capacity limit as a result of purely theoretical interest.

For example, voice-grade telephone channels can be quite accurately modeled as bandlimited Gaussian channels. The capacity of typical channels exceeds 20,000 b/s, and the commercial incentive for achieving the highest possible data rates is high. Nonetheless, 20 years after Shannon, the upper limit on practically achievable data rates was about 2,400 b/s. Thirty years after Shannon, the upper practical limit was generally considered to be 9,600 b/s. Furthermore, the advance to 9,600 b/s had essentially nothing to do with Shannon's work, but rather was due primarily to the implementation of digital adaptive equalization.

In the past decade, a series of advances has occurred that now make it possible to assert that channel capacity can be approached quite closely, not just in theory but also in certain practical applications (such as telephone-line modems).

The most striking of these advances was the invention of trellis-coded modulation (TCM) by Ungerboeck [3]. On ideal Gaussian channels (white Gaussian noise, no intersymbol interference), for symbol error rates of the order of 10^{-6}, Ungerboeck's schemes can achieve coding gains of from 3 dB to almost 6 dB, and can thus close most of the gap of approximately 9 dB between the Shannon limit and the rate achievable with uncoded modulation. (For a review of trellis-coded modulation, see Ungerboeck [4].)

In the past few years, it has been recognized that even more significant gains can be achieved by increasing the signaling rate so as to use the maximum possible channel bandwidth, and then using powerful equalization techniques to cope with the severe distortion that typically results from such bandwidth expansion. Naturally, it is desirable to do this in combination with coded modulation.

This paper reviews recently developed techniques that achieve these joint objectives. These techniques generalize the precoding techniques for uncoded modulation that were proposed independently more than 20 years ago by Tomlinson [5] and by Harashima and Miyakawa [6, 7]. Practically, these techniques have been used in telephone-line modems to achieve data rates of 19.2 kb/s and above, and are now being considered for international standardization. Theoretically, it has been shown that with these techniques, the capacity of any strictly bandlimited, high-SNR Gaussian channel can be approached as closely as can the capacity of an ideal Gaussian channel.

We begin with a review of Shannon's results and then give a brief history of telephone-line modems. After describing the principal classes of classical equalization techniques (linear equalization [LE], decision-feedback equalization [DFE], and maximum-likelihood sequence estimation [MLSE]) we introduce Tomlinson-Harashima precoding. We show that precoding combines nicely with coded modulation and with a technique known as shaping. Finally, we show that, in principle, optimized combinations of coding, shaping, and precoding can achieve nearly the channel capacity of any strictly bandlimited, high-SNR Gaussian channel, and in practice can approach capacity within about 3 dB or less. We conclude that the precoding approach is a practical route to capacity on high-SNR, bandlimited channels (e.g., telephone lines), and that the decision-feedback equalization structure that it embodies is in some sense canonical for coded modulation.

G. David Forney, Jr. is Vice President of the technical staff at Motorola Inc., Codex.

M. Vedat Eyuboğlu is manager of transmission research and advanced development at Codex Corporation, a subsidiary of Motorola, Inc.

[1] The precise quote is: "It is unfortunately true that the search for good codes with large [size] |S| has thus far been unrewarding. However, as we have seen, almost all codes are good. Thus we are tempted to infer that any code of which we cannot think is good."

 0163-6804/91/$01.00 1991© IEEE

Voice-grade telephone channels can be quite accurately modeled as bandlimited Gaussian channels.

Capacity of Gaussian Channels

A *Gaussian channel* is a combination of a linear filter and additive Gaussian noise, illustrated in Fig. 1. In the time domain, the output of such a channel is given by

$$r(t) = x(t) \cdot h(t) + n(t),$$

where $x(t)$ is the transmitted signal, $h(t)$ is the channel impulse response, $x(t) \cdot h(t)$ is the convolution of $x(t)$ with $h(t)$, and $n(t)$ is zero-mean colored Gaussian noise. In the frequency domain, the channel may be characterized by its frequency response $H(f)$, the Fourier transform of $h(t)$, and the signal and noise may be characterized by their power spectra $S(f)$ and $N(f)$. The total transmit power is limited by an average power constraint, $\int S(f)df \le P$.

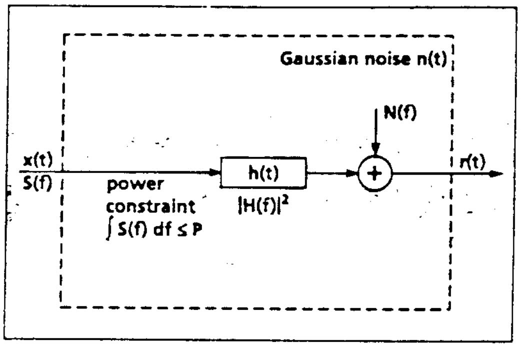

■ **Figure 1.** *Gaussian channel model*

The *channel SNR function* $|H(f)|^2/N(f)$ summarizes the channel characteristics for most purposes.

An ideal bandlimited Gaussian channel is characterized by a brickwall linear filter, shown schematically in Fig. 2, for which $H(f)$ is equal to a constant over some frequency band B of width W Hz and equal to zero elsewhere, and by white Gaussian noise, whose power spectrum $N(f)$ is equal to a constant $N_0/2$ over B. More generally, it is sufficient that the channel SNR function $|H(f)|^2/N(f)$ be constant over B and zero elsewhere.

Quadrature amplitude modulation (QAM) may be used over such a channel to send a sequence x of two-dimensional symbols x_k from a discrete signal set at a rate of W symbols per second. With perfect filtering and receiver synchronization, the sampled channel outputs r_k are given by

$$r_k = x_k + n_k,$$

■ **Figure 2.** *Ideal bandlimited Gaussian channel*

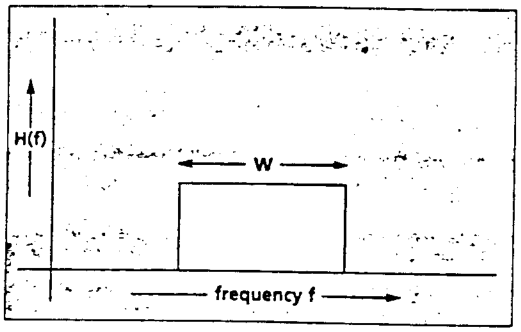

in complex notation, where x_k is the corresponding transmitted symbol, and n_k is a noise sample from a discrete-time white Gaussian noise process. In other words, there is no intersymbol interference (ISI), and noise samples are statistically independent. The average energy per two-dimensional symbol of the signal and noise samples will be denoted by S_x and S_n, respectively; for the ideal channel, $S_x = P/W$ and $S_n = N_0$.

Under a transmit average power constraint, the channel capacity of an ideal Gaussian channel is given by

$$C = \log_2 (1 + S_x/S_n)$$

bits per two-dimensional symbol (bits per second per Hz), or

$$\bar{C} = CW = W \log_2 (1 + S_x/S_n)$$

bits per second, where $S_x/S_n = P/N_0W$. An ideal Gaussian channel can therefore be characterized by two parameters: its bandwidth W and its signal-to-noise ratio (SNR) S_x/S_n, or equivalently its capacity C in b/s/Hz.

For example, a telephone channel typically has an SNR of approximately 28 dB to 36 dB or more, and a bandwidth of 2,400 Hz to 3,200 Hz or more. Although telephone channels are by no means ideal, their capacity can be crudely estimated by this formula to be of the order of 9 bits per Hz to 12 bits per Hz, or from 20,000 b/s to more than 30,000 b/s.

Capacity of Non-Ideal Channels

The capacity of an arbitrary Gaussian channel may be determined as follows. Let $S(f)$ be an arbitrary transmit power spectrum over a transmission band B. Divide B into many narrow subbands of width Δf, where Δf is chosen small enough so that the channel SNR function $|H(f)|^2/N(f)$ is approximately constant within any one subband. Then, given $S(f)$, the capacity within a subband of width Δf around frequency f_i is

$$C(f_i) = \log_2 [1 + S(f_i)|H(f_i)|^2/N(f_i)]$$

bits per symbol, or $C(f_i)\Delta f$ bits per second. The aggregate capacity across the whole band is given by

$$\bar{C} = \sum_i C(f_i)\Delta f \to \int_B C(f)\, df = \int_B \log_2 [1 + S(f)|H(f)|^2 / N(f)]\, df$$

bits per second.

The optimum transmit power spectrum $S(f)$ may be determined by water-pouring, as shown in Gallager [8], and illustrated in Fig. 3. The optimum transmit spectrum $S(f)$ is nonzero over a band of frequencies $B = \{f: N(f)/|H(f)|^2 < K_0\}$, called the capacity-achieving band, and within B is equal to

$$S(f) = K_0 - N(f)/|H(f)|^2,$$

where K_0 is a threshold chosen so that the transmit power spectrum $S(f)$ satisfies the power constraint $\int_B S(f)\, df \le P$.

For the ideal brickwall channel of Fig. 2, for example, the optimum transmit spectrum is constant over the band where $|H(f)|^2/N(f) > 0$ and zero elsewhere, and the capacity expression reduces, as before, to

$$\bar{C} = W \log_2 (1 + S_x / S_n)$$

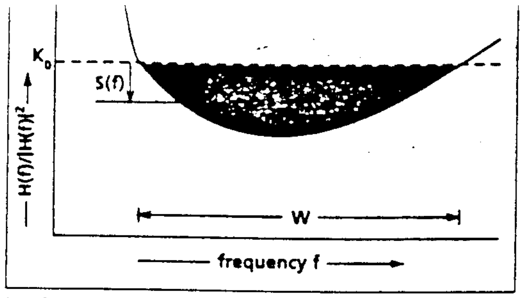

■ Figure 3. *Determination of optimum water-pouring spectrum*

The channel power spectrum $|H(f)|^2$ of typical telephone channels is reasonably constant over the center of the voice band, but drops sharply at the band edges. The location of the band edges depends on the carrier facilities used in the connection; the lower band edge is typically in the range between 100 Hz and 500 Hz; the upper band edge is between 3,000 Hz and 3,600 Hz. On such channels, calculations show that it is important to optimize the transmission band B, but that typically a flat transmit spectrum is almost as good as the optimum water-pouring spectrum, if B is nearly optimal [9].

Adaptive-Bandwidth, Adaptive-Rate Modulation

The approach to optimum use of nonideal channels that will be discussed in this paper is as follows. Using an initial probing signal, the receiver measures the channel characteristics and sends them to the transmitter via a reverse channel. The transmitter then determines how best to signal, including finding a band B that approximates the optimum transmission band. If B has width W, then it uses QAM modulation at a rate of W symbols per second, with an appropriate carrier frequency. It may also optimize the transmit power spectrum; as indicated in the previous paragraph, a flat spectrum usually is near-optimal, so typically this step is unnecessary. It also chooses a coded modulation scheme with rate R bits per symbol (b/s/Hz) as high as possible, given the desired error probability; that is, a rate R as close as possible to

$$C = \bar{C} / W = [\int_B \log_2[1 + S(f)|H(f)|^2 / N(f)]df] / W$$

The transmitter and receiver cooperate to compensate for (equalize) the channel SNR function $|H(f)|^2/N(f)$, which otherwise would cause significant intersymbol interference (ISI) and/or noise correlation.

Such an adaptive-bandwidth, adaptive-rate scheme is suitable for the point-to-point two-way applications in which telephone-line modems are used, and to which the capacity theorem is directed. In order to approach capacity, the transmitter needs to know the channel characteristics. This may not be possible on one-way, broadcast, or rapidly time-varying channels, unless all channel characteristics are known quite well *a priori*.

An alternative adaptive-bandwidth, adaptive-rate method of approaching the channel capacity of arbitrary Gaussian channels is multicarrier modulation [10]. As in the foregoing argument, multicarrier modulation divides the band into independent narrow subbands of width Δf, and each subband is independently modulated at a symbol rate of $1/\Delta f$ with power S(f). If the subbands are narrow enough, and a small guard interval is provided between symbols, it can be assumed that intersymbol interference and crosschannel interference are negligible; i.e., ordinary equalization is not required. For a given transmit power distribution S(f) across subbands, the capacity of a subband is approximately equal to

$$C(f) = \log_2 [1 + S(f)|H(f)|^2/N(f)]$$

bits per symbol, or $C(f)\Delta f$ bits per second, which can in principle be approached with powerful coded modulation.

Multicarrier modulation is in principle a general method of solving the combined equalization and coding problem. In practice, its inherent delay due to its use of a long symbol interval $1/\Delta f$ rules out its use for some modem applications, although it is suitable for others. Some work has been done on multicarrier modulation with few carriers and short symbol intervals [11]; then intersymbol interference arises and must be equalized.

History of Equalization in Telephone-Line Modems

The progress of equalization technology probably has been the most important factor in increasing the achievable data rates of telephone-line modems.

Table I gives a brief history of the CCITT Recommendations that have been developed over time as international standards for telephone-line modems. Most of these Recommendations (V.32 and V.33 being the exceptions) were based on earlier commercial modems. (For an earlier version of Table I focusing on modulation rather than equalization, see [12]. For a more complete historical survey, see [13].)

The earliest synchronous modem was the Bell 201 (c. 1962), standardized in 1968 in CCITT Recommendation V.26. This modem used four-phase (R = 2) modulation at up to 1,200 symbols per

The progress of equalization technology probably has been the most important factor in increasing the achievable data rates of telephone-line modems.

■ Table 1. *Modem milestones*

Date	Rec.	Rate(b/s)	R(b/s/Hz)	W(Hz)	Modulation	Equalization
1968	V.26	2400	2	1200	4-phase	fixed
1972	V.27	4800	3	1600	8-phase	manual
1976	V.29	9600	4	2400	16-QAM	adaptive, linear
1986	V.32	9600	4	2400	2D TCM	adaptive, linear
1986	V.33	14400	6	2400	2D TCM	adaptive, linear
1993?	'V.fast'	2[illegible]000?	up to 8?	up to 3200?	4D TCM?	precoding?

The 1960s were a time of consider-able research into adaptive equalization, notably that of Lucky and his colleagues at AT&T Bell Labs.

second to send up to 2400 b/s in a nominal W = 1200 Hz band. It used a fixed equalizer, designed for typical lines.

About 1967, the Milgo 4400/48 successfully achieved a rate of 4800 b/s by extending the usable band to W = 1600 Hz, and using eight-phase (R = 3) modulation at 1,600 symbols per second. This modem incorporated a manually adjustable equalizer to equalize the wider band; the operator adjusted a knob on the front panel to zero a null meter. (The imagination boggles at the idea of today's personal computer or facsimile user trying to twiddle a knob on his modem to obtain good data, similar to a listener trying to pull in a distant night game on an old radio set.) CCITT Recommendation V.27 in 1972 was based on this modem.

The 1960s were a time of considerable research into adaptive equalization, notably that of Lucky and his colleagues at AT&T Bell Labs [14]. The implementation technology of that time was primitive, and much of this work focused on adaptation algorithms that did not require multiplications (e.g., tapped delay lines with a zero-forcing algorithm which could be implemented by up/down counters).

In 1971, the Codex 9600C successfully achieved a rate of 9,600 b/s over private lines, using an automatic adaptive digital linear equalizer to equalize a W = 2400 Hz band, and a nonstandard 16-point QAM signal set (R = 4) at 2,400 symbols per second to combat combined noise and phase jitter. CCITT Recommendation V.29 was based on this modem.

During the 1970s, modems became smaller, cheaper, more reliable, and more versatile, but not faster. (For a survey of advances in digital signal processing VLSI from that period to the present, see [15].) The principal practical advances in equalization were in fast-training algorithms for multipoint and half-duplex applications, and in the development of fractionally spaced equalizers. (For a survey of adaptive equalization theory and practice in this period, see Qureshi [16].) Even when the first 14,400 b/s modems appeared (the Paradyne 14400, c. 1981), they still used uncoded modulation, a fixed bandwidth of 2400 Hz and linear equalization.

During 1983-1984, Recommendation V.32 for 9600 b/s transmission over ordinary dial lines was developed, just after the publication of Ungerboeck's landmark paper introducing TCM [3]. It was quickly recognized that the 3-4 dB of coding gain provided by simple TCM schemes was just what was needed for robust transmission at 9600 b/s over a wide variety of dial lines. Recommendation V.32 embodied an eight-state two-dimensional (2D) nonlinear code due to Wei [17], similar to Ungerboeck's eight-state 2D code, but with a clever new solution to the problem of rotational invariance.

Another key advance in V.32 was the incorporation of adaptive echo cancellation for noninterfering two-way transmission over a single two-wire circuit [18]. Otherwise, this standard used R = 4 b/s/Hz transmission over a fixed 2400 Hz bandwidth and linear equalization, continuing previous practice.

Recommendation V.33 applied the same TCM scheme with larger signal sets to support R = 6 b/s/Hz and thus a data rate of 14.4 kb/s over private lines. Recently V.32 has been upgraded to achieve 14.4 kb/s over dial lines (V.32bis) in the same manner.

In 1985, the Codex 2680 successfully achievec a rate of 19.2 kb/s over private lines by using more powerful multidimensional TCM schemes also due to Wei [19, 20] to support R = 7 b/s/H at W = 2,743 symbols per second. The use of a wider bandwidth was a significant advance; as indicated by the capacity formula, 1 b/s/Hz equates to 3 dB in SNR at high rates, so achieving 19.2 kb/s with W = 2400 and R = 8 would have required about 2 dB more margin (3 dB, minus approximatel 1 dB due to the expansion of the noise bandwidth by a factor of 8/7). The wider bandwidth wa equalized with powerful but still linear equalization techniques.

During this period, Telebit introduced ar adaptive-rate, adaptive-bandwidth modem of the multicarrier type, primarily for half-duplex applications on dial lines. This modem, which supported data rates up to about 20 kb/s, has probabl been the most successful multicarrier modem ever offered commercially, and has drawn increase attention to the multicarrier approach.

In 1990, the Codex 3600 successfully achieve rates of up to 24 kb/s on private (four-wire) line by using the adaptive-bandwidth, adaptive-rat approach described above. Combined equalization multidimensional TCM and shaping were achieve using trellis precoding (to be described later) Similar techniques are now being considered fo standardization in the forthcoming "V.fast" Recommendation for dial (two-wire) modems a rates of 19.2 kb/s and higher.

Classical Equalization Techniques

This section discusses the classical equalizatio techniques designed for uncoded pulse amplitude modulation (PAM) or uncoded QAM: linear equalization (LE), decision-feedback equalizatio (DFE), and maximum likelihood sequence estimation (MLSE).

Without loss of optimality, subject to certain technical conditions, the front end of a receiver in a QAM modem may be taken to consist of a whitene matched filter (WMF) and a symbol-rate sample [21]. Thus, the equivalent discrete-time channe shown in Fig. 4 is obtained. The received sequenc $\{r_k\}$ is given by

$$r_k = \sum_{j\geq 0} h_j x_{k-j} + n_k = y_k + n_k,$$

where $\{x_k\}$ is the transmitted sequence, $\{h_j\}$

■ **Figure 4.** *Equivalent discrete-time channel mo using WMF*

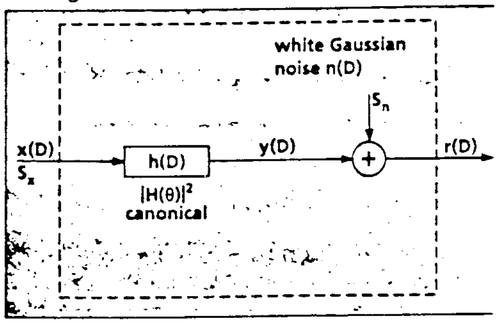

the sequence of coefficients of an equivalent discrete-time channel response, and $\{n_k\}$ is a discrete-time white Gaussian noise sequence. The term $y_k = \sum_{j\geq 0} h_j x_{k-j}$ is the noiseless output. (In a QAM system, all of these quantities are complex numbers, but the reader may equally well think in terms of a one-dimensional PAM system in which all quantities are real.) In D-transform notation,

$$r(D) = x(D)h(D) + n(D) = y(D) + n(D).$$

The equivalent discrete-time channel response h(D) is *canonical*, i.e., causal ($h_j = 0$ if $j < 0$), monic ($h_0 = 1$), and minimum-phase (all poles are outside the unit circle, and all zeroes are on or outside the unit circle). Its frequency response is $H(\theta)$, $-\pi < \theta \leq \pi$. The channel is ideal if $h(D) = 1$, or equivalently if $H(\theta) = 1$, $-\pi < \theta \leq \pi$. The energy of the response is

$$\|h\|^2 = \sum_j |h_j|^2 = (1/2\pi)\int_{-\pi}^{\pi}|H(\theta)|^2\, d\theta \ ;$$

$\|h\|^2$ is finite, and $\|h\|^2 \geq 1$, with equality if and only if the channel is ideal, $h(D) = 1$.

Both x(D) and n(D) are white (iid) sequences, with average energies S_x and S_n per symbol, respectively.

Zero-forcing linear equalization (ZF-LE) may be used if the channel response h(D) has a well-defined inverse 1/h(D). With ZF-LE, the received sequence r(D) is simply filtered by 1/h(D) to produce an equalized response

$$r'(D) = r(D)/h(D) = x(D) + n(D)/h(D) = x(D) + n'(D),$$

as shown in Fig. 5. Symbol-by-symbol decisions can then be made on the transmitted sequence x(D). Intersymbol interference is eliminated, but the average noise energy is enhanced by the energy $\|1/h\|^2$ of the inverse filter response 1/h(D):

$$S_{n'} = S_n\|1/h\|^2 = S_n[(1/2\pi)\int_{-\pi}^{\pi}(1/|H(\theta)|^2)\, d\theta].$$

The signal-to-noise ratio at the decision point is thus equal to

$$SNR_{ZF\text{-}LE} = (S_x/S_n)[(1/2\pi)\int_{-\pi}^{\pi}(1/|H(\theta)|^2)\,d\theta]^{-1}.$$

The noise enhancement factor $\|1/h\|^2$ is greater than or equal to 1, with equality if and only if the channel is ideal, $h(D) = 1$. If the response $|H(\theta)|^2$ is reasonably constant over the Nyquist band $\{-\pi < \theta \leq \pi\}$, then the noise enhancement of ZF-LE is not very serious. However, if $|H(\theta)|^2$ has a near-null, then noise enhancement can become large. If $|H(\theta)|^2$ has an actual null, then h(D) is not invertible and the ZF-LE structure is not well defined.

This analysis indicates that linear equalizers can be satisfactory in a QAM modem whenever the channel has no nulls or near-nulls within the nominal Nyquist band. In order to approach capacity, however, the transmission band must be expanded to the entire usable bandwidth of the channel, which inevitably leads to severe attenuation at the band edges; then linear equalization no longer suffices.

The most popular nonlinear equalizer structure is decision-feedback equalization (DFE). *Zero-forcing decision-feedback equalization* (ZF-DFE) is illustrated in Fig. 6. The basic idea, when estimating x_k, is to assume that all previous estimates are correct: $\hat{x}_{k-j} = x_{k-j}, j \geq 1$ (the ideal DFE assumption). Then the "tail" $\sum_{j\geq 1} h_j x_{k-j}$ can be removed by subtraction; i.e., the equalized sample is

$$r_k' = r_k - \sum_{j\geq 1} h_j \hat{x}_{k-j} = x_k + n_k,$$

where the second equality depends on the ideal DFE assumption and the fact that $h_0 = 1$. In D-transform notation,

$$r'(D) = r(D) - \hat{x}(D)[h(D) - 1] = x(D) + n(D).$$

Thus, intersymbol interference is completely removed, and the noise is white. The signal-to-noise ratio is simply equal to

$$SNR_{ZF\text{-}DFE} = S_x/S_n,$$

which improves on $SNR_{ZF\text{-}LE}$ by the noise enhancement factor $\|1/h\|^2$. Thus, $SNR_{ZF\text{-}DFE} \geq SNR_{ZF\text{-}LE}$, with equality if and only if $h(D) = 1$.

The optimum equalization structure for a discrete transmitted sequence x(D) in the presence of ISI generally is considered to be maximum-likelihood sequence estimation (MLSE). If the transmitted symbols x_k are drawn from an M-point signal set, and the channel response h(D) has length υ (or closely approximates such a finite response), then the channel can be modeled as a shift register of length υ with M-state memory elements; i.e., as an M^{υ}-state machine. An M^{υ}-state Viterbi algorithm (VA) may be used to implement MLSE for such a finite-state system [21]. If M and/or υ is even moderately large, however, an M^{υ}-state VA generally is considered to be too complex to implement.

On many channels without severe ISI, an MLSE detector achieves the effective SNR of the *matched filter bound:*

$$SNR_{MFB} = S_x\|h\|^2 / S_n.$$

The matched filter bound often is taken as a bound

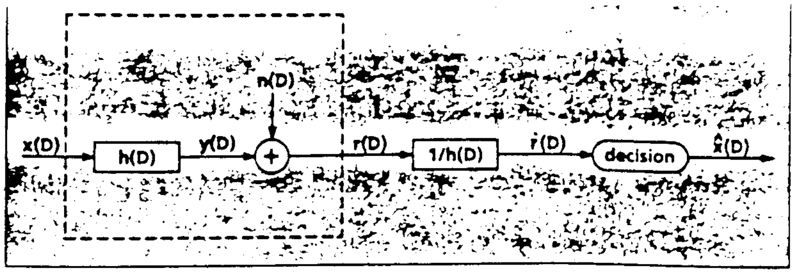

■ **Figure 5.** *Zero-forcing linear equalization*

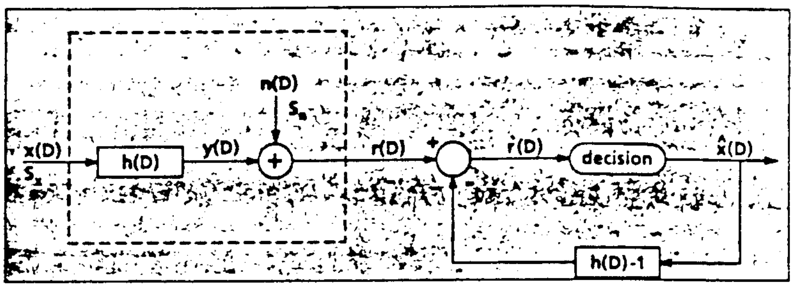

■ **Figure 6.** *Zero-forcing decision-feedback equalization*

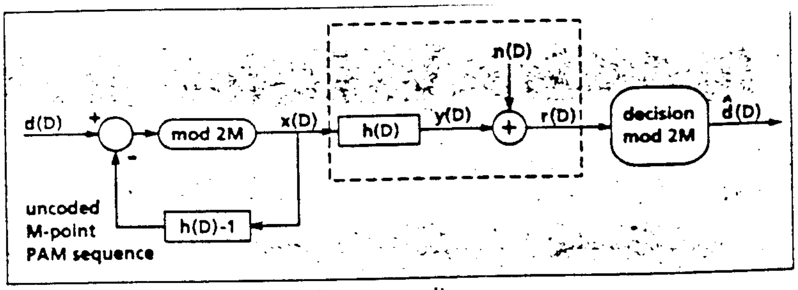

■ Figure 7. *Tomlinson-Harashima precoding*

on the best possible SNR achievable on a channel with response h(D). When ISI is severe, however, even MLSE can fail to achieve this SNR, and performance analysis becomes difficult.

Tomlinson-Harashima Precoding

Tomlinson-Harashima precoding was invented independently and more or less simultaneously (1968-1969) in theses by Tomlinson in the United Kingdom [5] and Harashima in Japan [6, 7]. In the English-language literature it generally has been known as Tomlinson precoding, but here it will be called by the names of both inventors. If no confusion can occur, we shall simply call it precoding.

Like the equalization techniques in the previous section, Tomlinson-Harashima precoding was proposed for use with uncoded modulation. Unlike those techniques, in which all equalization takes place in the receiver, precoding is a transmitter technique.

Precoding is more closely tied to the specific signal set than the previous equalization techniques, although as we shall see it can be generalized to all cases of interest. Originally it was proposed for use with M-point PAM, where the signal set A consists of M equally spaced levels, e.g.,

$$A = \{\pm 1, \pm 3, \ldots, \pm (M-1)\},$$

the set of all odd integers in the interval [-M, M].

Tomlinson-Harashima precoding works as follows, as illustrated in Fig. 7. The equivalent discrete-time channel response h(D) is assumed to be known at the transmitter. (An adaptive linear equalizer in the receiver can be used to assure that the channel response remains equal to a preset h(D), if the channel does not vary greatly.) The transmitter generates a data sequence d(D) whose symbols d_k are in the PAM signal set A. The transmitted signal x_k is formed by first subtracting the "tail" $\sum_{j\geq 1} h_j x_{k-j}$ due to previously transmitted signals from d_k (decision feedback in the transmitter), and then reducing $d_k - \sum_{j\geq 1} h_j x_{k-j}$ modulo 2M to the half-open interval (-M, M].

The modulo 2M operation can be characterized in various ways. By definition, the transmitted symbol x_k is the unique number that satisfies the two constraints

$$x_k \equiv d_k - \sum_{j\geq 1} h_j x_{k-j} \text{ modulo } 2M;$$

$$x_k \in (-M, M].$$

In other words, the transmitter finds the unique integer z_k such that

$$x_k = d_k + 2Mz_k - \sum_{j\geq 1} h_j x_{k-j}$$

is in the interval (-M, M]. In D-transform notation,

$$x(D) = d(D) + 2Mz(D) - x(D)[h(D) - 1].$$

This reduces to

$$x(D) = [d(D) + 2Mz(D)]/h(D) = y(D)/h(D),$$

where y(D) is a sequence of odd-integer modified data symbols $y_k = d_k + 2Mz_k$. Consequently, the received signal is

$$r(D) = x(D)h(D) + n(D) = y(D) + n(D).$$

Thus, in the absence of noise the received sequence is the modified data sequence y(D) = d(D) + 2Mz(D), which may be detected on a symbol-by-symbol basis to give an estimated sequence $\hat{y}(D)$. There is no ISI, and the error probability is the same as if the original data sequence d(D) were sent on the same channel and detected with a ZF-DFE, or if d(D) were sent on an ideal channel with signal-to-noise ratio $SNR_{ZF\text{-}DFE} = S_x/S_n$. An estimate $\hat{d}_k$ of the original data symbol d_k can then be retrieved by reducing $\hat{y}_k$ to the interval (-M, M] with a modulo 2M operation.

By construction, the transmitted symbol x_k lies in the interval (-M, M]. For a general channel response h(D), it will be randomly distributed over this continuous interval. (If h(D) is integer-valued, then x(D) will still be odd-integer-valued.) It is easy to show that if d(D) is precisely an iid sequence uniformly distributed over (-M, M], then so is x(D), regardless of h(D). If M is large, then d(D) statistically approximates a continuous iid uniform sequence; therefore x(D) also statistically approximates a continuous iid uniform sequence.

The average energy per symbol of a PAM data sequence d(D), assuming that the M points in the signal set A are equiprobable, is

$$S_x = (M^2 - 1)d_{min}^2 / 12 = (M^2 - 1)/3$$

The average energy of a continuous iid sequence uniformly distributed on (-M, M] is

$$\hat{S}_x = M^2/3$$

this is a good approximation to the average power of the transmitted sequence x(D) when M is large. The power increases by a factor of $M^2/(M^2 - 1)$, which is insignificant when M is large.

For large M, the transmitted sequence x(D) is approximately iid, so the transmit power spectrum is approximately flat, which is close to the optimal water-pouring spectrum on high-SNR channels. If spectral shaping is desired, then x(D) may be filtered, just like a PAM signal [22].

The set of modified symbols y_k that can actually occur for a given response h(D) is in general an expanded odd-integer signal set of the PAM type. The received signal range is correspondingly expanded. This effect is not troublesome; statistically, the output of the channel is approximately the same as if an ordinary PAM signal were sent.

Note that precoding works even if h(D) is not invertible; i.e., if $\|1 / h\|^2$ is infinite.

The key points about Tomlinson-Harashima precoding may be summarized as follows:

- The transmitter is assumed to know the channel response h(D).
- The channel output sequence y(D) = d(D) + 2Mz(D) is chosen so that the symbols in the transmitted sequence x(D) = y(D)/h(D) lie in the interval (-M, M].
- For large M, the transmitted sequence x(D) statistically resembles an ordinary PAM sequence, except that its possible values are in general continuous in the interval (-M, M].
- The receiver operates in essentially the same way as it would for ordinary PAM on an ideal channel (i.e., it performs symbol-by-symbol detection).
- Error probability is the same as with ideal ZF-DFE; i.e., the same as on an ideal channel with signal-to-noise ratio $SNR_{ZF\text{-}DFE} = S_x/S_n$.

Apart from the first point, Tomlinson-Harashima precoding appears to be an attractive alternative to ZF-DFE. In fact, with uncoded systems, it has not received much attention in practice. Its performance is no better than that of ZF-DFE under the ideal DFE assumption, and with uncoded systems the ideal DFE assumption usually is a good one, since the feared phenomenon of catastrophic DFE error propagation rarely materializes. Therefore, DFE has generally been preferred to precoding, since it does not require information about the channel at the transmitter.

Combined Precoding and Coded Modulation

Now suppose that it is desired to use nonlinear equalization in combination with trellis-coded modulation. With DFE, a major problem then arises: symbol-by-symbol decisions are unreliable, but reliable decoded decisions are not available until after a decoding delay.

One approach is to use an interleaver to shuffle the symbols in such a way as to provide reliable decisions for feedback most of the time [23]. Such an interleaver, however, must have a large span in order for most feedback decisions to be reliable, and thus may introduce considerable delay.

Another approach that is optimum in principle is to consider the combination of the trellis encoder and the channel as one large finite-state machine, whose state space is the product of the encoder and channel state spaces, and then to perform MLSE via the Viterbi algorithm on the entire system. Generally, however, this is much too complicated.

Various reduced-state sequence estimation (RSSE) algorithms have been investigated, with the intent of significantly reducing the number of states in the above MLSE approach without

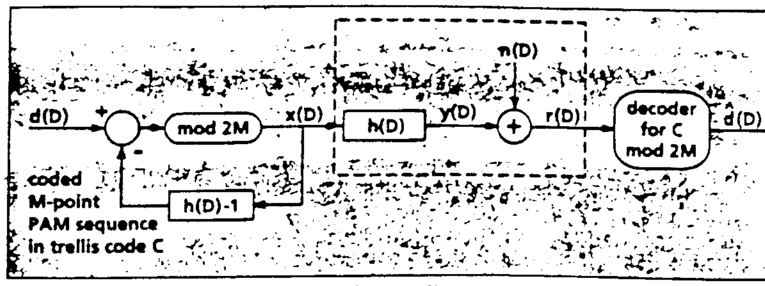

Figure 8. *Combined coding and precoding*

too much loss of performance [24-27]. The simplest and perhaps most attractive RSSE scheme, called "parallel decision feedback decoding," simply uses the original encoder trellis, but performs decision-feedback equalization on each surviving path in the VA individually, based on the history of that path [25-27]. The performance and complexity of RSSE schemes seems to depend somewhat unpredictably on how the code structure and channel response interact. Nonetheless, RSSE may be attractive, particularly when equalization must be performed wholly in the receiver.

Often when the transmitter and receiver can cooperate, precoding is a more attractive approach. Precoding can be combined with coded modulation essentially with "no glue," as follows.

Practically all known good trellis codes work by constraining the one or two least significant bits in a binary representation of the data symbols d_k according to the possible code sequences in a binary convolutional code (depending on whether C is a mod-2 or a mod-4 code [20]). In other words, a PAM data sequence d(D) is a legitimate word in a given trellis code C if, and only if, the least significant bits of the components d_k of d(D) satisfy certain constraints.

The key point is that if d(D) is a coded PAM sequence in such a trellis code C, if z(D) is any integer sequence, and if M is a multiple of 4, then the modified data sequence

$$y(D) = d(D) + 2Mz(D)$$

also is a sequence in C.

A combined coder-precoder can therefore operate as shown in Fig. 8. The transmitter first generates a coded PAM sequence d(D) in a given trellis code C. It then precodes it as in the previous section to generate a transmitted sequence x(D) = y(D)/h(D), where y(D) = d(D) + 2Mz(D). If M is a multiple of 4, then y(D) also is in C.

The received sequence is then r(D) = y(D) + n(D); i.e., a code sequence in C plus white Gaussian noise. An ordinary Viterbi algorithm can operate as usual on r(D) to find the closest code sequence $\hat{y}(D)$ to r(D), which can then be reduced modulo 2M to $\hat{y}(D)$. For large M, the performance and complexity of such a decoder will be essentially the same as if coded sequences d(D) were sent through an ideal Gaussian channel with signal-to-noise ratio $SNR_{ZF\text{-}DFE} = S_x/S_n$ and detected with a VA decoder for C.

In summary, precoding can be combined with coding in a very natural way. The combined system obtains the equalization gain of ideal DFE in combination with the coding gain of the given trellis code C.

At high rates, shaping gain is essentially independent of coding gain, and the two problems can be addressed separately.

Trellis Precoding: Combined Precoding, Coding, and Shaping

In the combined coding-precoding system of Fig. 8, as in uncoded Tomlinson-Harashima precoding, the transmitted symbols are essentially uniformly distributed within the interval (-M, M]. This makes it impossible to achieve any shaping gain.

Briefly, *shaping gain* is a reduction in average signal energy that can be obtained by shaping N-dimensional signal sets more like an N-sphere than an N-cube, or equivalently by causing the probability distribution of one- or two-dimensional transmitted signals to be more like a Gaussian distribution than a uniform distribution. The maximum possible shaping gain is $\pi e/6$, or 1.53 dB. This gain is obtained if and only if the transmitted sequence is statistically equivalent to a white Gaussian sequence. At high rates, shaping gain is essentially independent of coding gain, and the two problems can be addressed separately [28].

While shaping gain is limited to 1.53 dB, and therefore is much less than typical coding gains, it nonetheless is worth pursuing. On ideal channels, simple techniques have been developed that achieve of the order of 1 dB of shaping gain with modest complexity [29, 30]. Given a trellis code with 3 dB or 4 dB of coding gain, it is much easier to obtain the next 1 dB with such shaping techniques than by more complicated coding.

A generalization of Tomlinson-Harashima precoding, called trellis precoding, can be used to achieve the coding gain of a given trellis code, the equalization performance of ideal ZF-DFE, and also shaping gain close to the ultimate shaping gain of 1.53 dB [31].

In brief summary, trellis precoding is a further generalization of the combined coding-precoding system illustrated in Fig. 7. The channel output sequence is again a modified data sequence; i.e., y(D) = d(D) + 2Mz(D). However, instead of choosing z(D) from all integer-valued sequences, the transmitter searches for the best integer sequence z(D) in a given trellis code C_s, called the *shaping code*. The signal alphabet for the data sequence d(D) must be chosen appropriately so that d(D) can be uniquely recovered from the channel output sequence y(D).

The search criterion usually is to choose the z(D) in C_s that minimizes the average energy of the transmitted sequence x(D) = y(D)/h(D). However, other criteria, such as peak constraints on the transmitted symbols x_k, may also be used.

It is shown in [31] that the best x(D) may be found by a simple RSSE-like search through the trellis of the shaping code C_s. Even with a simple four-state code C_s, shaping gains (average signal energy reductions) approaching 1 dB may be obtained. The resulting transmitted sequence x(D) begins to statistically resemble a peak-limited white Gaussian sequence more than an iid uniform sequence. With more complicated shaping codes, x(D) becomes still more Gaussian-like, and shaping gains approach the ultimate shaping gain of 1.53 dB.

Again, the least significant bits of d(D) are preserved in the output sequence y(D) = d(D) + 2Mz(D), so it follows that if d(D) is a sequence in a given trellis code C, then so is y(D). The received sequence r(D) = y(D) + n(D) may therefore again be decoded by an ordinary VA decoder for C. Such a decoder obtains the same coding gain as is obtained with the given code C on an ideal channel with signal-to-noise ratio $SNR_{ZF\text{-}DFE} = S_x/S_n$.

The key points about trellis precoding may be summarized as follows:

- The transmitter is assumed to know the channel response h(D).
- The noiseless channel output sequence y(D) = d(D) + Mz(D) is chosen to minimize the average energy of the transmitted signal sequence x(D) = y(D)/h(D), as well as to satisfy other practical criteria such as peak constraints.
- At high rates, the transmitted signal sequence x(D) statistically resembles a white Gaussian sequence, and thus achieves shaping gain.
- The receiver operates in essentially the same way as it would for the same trellis code on an ideal channel (i.e., it executes a standard Viterbi algorithm).
- Coding gain is the same as on an ideal channel with signal-to-noise ratio $SNR_{ZF\text{-}DFE} = S_x/S_n$. In other words, the full coding gain of the trellis code is obtained, in combination with the equalization performance of ideal ZF-DFE.

Trellis precoding has been implemented with symbol rates up to W = 3200 QAM symbols per second, in combination with multidimensional TCM. Test results [32] have shown that on lines on which V.32bis modems can achieve 14.4 kb/s (R = 6, W = 2400), a modem of this type can usually achieve of the order of 19.2 kb/s (R = 6, W = 3200). Furthermore, rates up to the order of 25.6 kb/s (R = 8, W = 3200) can be achieved on higher-quality lines.

Price's Result, and Attaining Channel Capacity

In a significant early paper, Price [33] observed that at high signal-to-noise ratios the dB gap between the performance of uncoded PAM with a ZF-DFE is essentially the same on all strictly bandlimited channels, regardless of the channel characteristics (provided that the ZF-DFE is well defined). He also pointed out that the same performance could alternatively be achieved with Tomlinson precoding. The importance of this result for trellis coding was noted in [23]. In this section we sketch Price's result and its consequences.

Consider an arbitrary Gaussian channel that is strictly bandlimited to some frequency band B of width W. Suppose that the transmitter uses uncoded QAM with an MxM square signal set at W symbols per second, with a transmit power spectrum that is flat over B and zero elsewhere, and that the receiver uses ideal ZF-DFE; or equivalently, that the system uses Tomlinson-Harashima precoding.

Define the *normalized signal-to-noise ratio* of an arbitrary transmission scheme with a data rate of R bits per two dimensions on this channel as

$$SNR_{norm} \triangleq SNR_{ZF-DFE} / (2^R - 1),$$

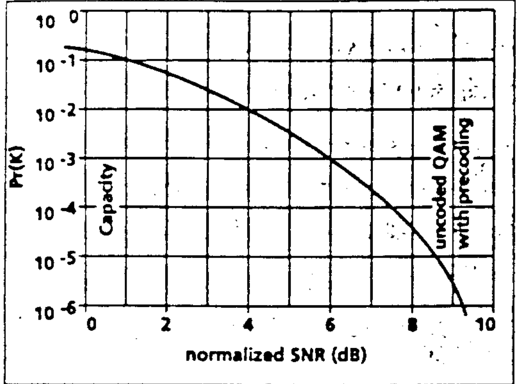

■ Figure 9. *Uncoded modulation performance vs. capacity on Gaussian channels*

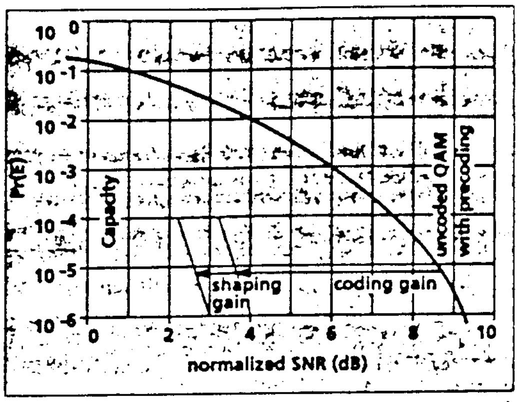

■ Figure 10. *Coded modulation performance with coding, shaping, and precoding*

It is not necessary to use the optimal MLSE decoding structure to operate near capacity.

where $SNR_{ZF\text{-}DFE} = S_x/S_n$. With an uncoded MxM square QAM signal set with minimum distance d_{min},

$$S_x = (M^2 - 1)d_{min}^2 / 6, \text{ and}$$

$$R = \log_2 M^2.$$

Since the noise is Gaussian, the symbol error probability can be well approximated at high rates by

$$\Pr(E) \cong 4 \bullet Q[(d_{min}^2 / 2S_n)^{1/2}].$$

$$= 4 \bullet Q[(3 \bullet SNR_{norm})^{1/2}],$$

where 4 is the number of near neighbors to interior points in the QAM signal set, and $Q(y) = \int_y^\infty p_{0,1}(x)dx$ is the Gaussian error probability function, with $p_{0,1}(x)$ a Gaussian distribution of mean 0 and variance 1. This gives a universal curve, independent of the data rate R, for the approximate performance of uncoded MxM QAM with ZF-DFE on arbitrary high-SNR, strictly bandlimited Gaussian channels, which is shown in Fig. 9.

According to Price's result, at high signal-to-noise ratios the capacity of this channel is well approximated by

$$C \cong \log_2 (1 + SNR_{ZF\text{-}DFE}),$$

in bits per two dimensional symbol. Therefore, the normalized signal-to-noise ratio can be approximated by

$$SNR_{norm} \cong (2^C - 1)/(2^R - 1).$$

Since C upperbounds the rate R per two dimensions of any feasible coded modulation scheme, the Shannon limit for an arbitrary strictly bandlimited, high-SNR Gaussian channel may be expressed in terms of SNR_{norm} by the simple bound

$$SNR_{norm} > 1.$$

Therefore, the Pr(E) curve of any feasible scheme must lie to the right of the vertical axis (0 dB) in Fig. 9, which has been marked "capacity." The gap between the Shannon limit and the performance of uncoded MxM QAM at $\Pr(E) \cong 10^{-6}$ may thus be seen to be about 9 dB. The dB gap between capacity and uncoded QAM is approximately the same for any strictly bandlimited, high-SNR Gaussian channel, including, of course, the ideal channel.

The main significance of Price's result is that, in combination with the results of previous sections, it shows that the dB gap between capacity and uncoded modulation performance can be closed to the same extent on arbitrary strictly bandlimited, high-SNR Gaussian channels as it can on ideal high-SNR channels. For example, suppose that there exists some coded modulation scheme with rate R, coding gain γ_c and shaping gain γ_s that achieves error probability

$$\Pr(E) \cong K \bullet Q[(3\gamma_c\gamma_s SNR_{norm})^{1/2}]$$

on an ideal channel at a certain SNR, where K is the error coefficient of the scheme and $SNR_{norm} = SNR_{ZF\text{-}DFE}/(2^R - 1)$. The performance that would be obtained with a coding gain of approximately 5 dB and a shaping gain of approximately 1 dB is illustrated in Fig. 10.

With precoding, we have seen that the same scheme can achieve the same coding gain γ_c with the same error coefficient K on an arbitrary channel characterized by the signal-to-noise ratio $SNR_{ZF\text{-}DFE}$. Furthermore, from empirical evidence, it appears that with trellis precoding it is possible to achieve a shaping gain γ_s as close to the ultimate shaping gain of $\pi e/6$ (1.53 dB) as desired. Therefore, the same improvement over the performance of uncoded systems can be obtained by a combination of coding, shaping, and precoding on arbitrary strictly bandlimited, high-SNR Gaussian channels as can be obtained by a combination of coding and shaping on ideal high-SNR channels.

The largest coding gains that have been demonstrated for Gaussian channels at error probabilities on the order of 10^{-6} are of the order of 5-6 dB; the largest shaping gains (on ideal channels), about 1.3 dB [30]. Therefore, it is possible to reduce the SNR gap of about 9 dB (at $\Pr(E) \cong 10^{-6}$) between capacity and uncoded modulation by approximately 7 dB on arbitrary strictly bandlimited, high-SNR Gaussian channels, using combined coding, shaping, and precoding. It is not necessary to use the optimal MLSE decoding structure to operate near capacity; the ZF-DFE structure will suffice (provided that it is well defined) with the feedback filter implemented in the transmitter via precoding.

A result of deBuda [34] shows that in principle lattice codes with spherical shaping can approach the channel capacity of an ideal high-SNR Gaussian channel arbitrarily closely. The above argument indicates that deBuda's result applies to arbitrary strictly bandlimited, high-SNR Gaussian channels.

Finally, it can be shown that these results can be generalized and extended by using minimum-mean-squared-error decision-feedback equalization (MMSE-DFE), rather than ZF-DFE [9]. The MMSE-DFE structure always is well defined, and its components may be used instead of ZF-DFE in a precoding system. There is a striking relationship between the signal-to-noise ratio of unbiased MMSE-DFE, $SNR_{MMSE\text{-}DFE,U}$, and the capacity C of an arbitrary Gaussian channel, namely

$$C = \log_2 (1 + SNR_{MMSE\text{-}DFE,U}).$$

This relationship, along with the other results of this paper, strongly suggests that MMSE-DFE is a canonical equalization structure for use with coded modulation.

References

[1] C. E. Shannon, "Communication in the Presence of Noise," *Proc. IRE*, vol. 37, pp. 10-21, January 1949.

[2] J. M. Wozencraft and B. Reiffen, *Sequential Decoding* (MIT Press, 1961).

[3] G. Ungerboeck, "Channel Coding with Multilevel/Phase Signals," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 55-67, 1982.

[4] G. Ungerboeck, "Trellis-Coded Modulation with Redundant Signal Sets," IEEE *Communications Magazine*, pp. 5-21, February 1987.

[5] M. Tomlinson, "New Automatic Equalizer Employing Modulo Arithmetic," *Electron. Lett.*, vol. 7, pp. 138-39, March 1971.

[6] H. Miyakawa and H. Harashima, "A Method of Code Conversion for a Digital Communication Channel with Intersymbol Interference," *Trans. Inst. Electron. Commun. Eng. Japan*, vol. 52-A, pp. 272-273, June 1969.

[7] H. Harashima and H. Miyakawa, "Matched-Transmission Technique for Channels with Intersymbol Interference," *IEEE Trans. Commun.*, vol. COM-20, pp. 774-80, August 1972.

[8] R. G. Gallager, *Information Theory and Reliable Communication* (Wiley, 1968).

[9] J. M. Cioffi, *et al.*, "MMSE Decision-Feedback Equalization and Coding—Part I: General Results," *IEEE Trans. Commun.*, to appear.

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

[11] J. Chow, J. Tu, and J. M. Cioffi, "A Discrete Multitone Transceiver System for HDSL Applications," *IEEE J. Select. Areas Commun.*, vol. SAC-9, pp. xxx-xxx, October 1991.

[12] G. D. Forney, *et al.*, "Efficient Modulation for Bandlimited Channels," *IEEE J. Select. Areas Commun.*, vol. SAC-2, pp. 632-47, September 1984.

[13] K. Pahlavan and J. L. Holsinger, "Voice-Band Data Communication Modems—A Historical Review: 1919-1988," *IEEE Communications Magazine* pp. 16-27, January 1988.

[14] R. W. Lucky, J. Salz, and E. J. Weldon, Jr., *Principles of Data Communication* (McGraw-Hill, 1968).

[15] H. M. Ahmed and R. B. Kline, "Recent Advances in DSP Systems," *IEEE Communications Magazine*, pp. 32-45, May 1991.

[16] S. U. Qureshi, "Adaptive Equalization," *Proc. IEEE*, vol. 73, pp. 1349-87, 1985.

[17] L.-F. Wei, "Rotationally Invariant Convolutional Channel Coding with Expanded Signal Space—Part II: Nonlinear Codes," *IEEE J. Select. Areas Commun.*, vol. SAC-2, pp. 672-686, September 1984.

[18] S. B. Weinstein, "A Passband Data-Driven Echo Canceller for Full-Duplex Transmission on Two-Wire Circuits," *IEEE Trans. Commun.*, vol. COM-25, pp. 654-66, January 1977.

[19] L.-F. Wei, "Trellis-Coded Modulation Using Multidimensional Constellations," *IEEE Trans. Inform. Theory*, vol. IT-33, pp. 483-501, July 1987.

[20] G. D. Forney, Jr., "Coset Codes—Part I: Introduction and Geometrical Classification," *IEEE Trans. Inform. Theory*, vol. IT-34, pp. 1123-51, Sept. 1988.

[21] G. D. Forney, Jr., "Maximum-Likelihood Sequence Estimation of Digital Sequences in the Presence of Intersymbol Interference," *IEEE Trans. Inform. Theory*, vol. IT-18, pp. 363-78, May 1972.

[22] A. R. Calderbank and J. E. Mazo, "Spectral Nulls and Coding with Large Alphabets," *IEEE Communications Magazine*, this issue, Dec. 1991. ADELE Please SUPPLY VOL., PAGE NO]

[23] M. V. Eyuboglu, "Detection of Coded Modulation Signals on Linear Severely Distorted Channels Using Decision-Feedback Noise Prediction with Interleaving," *IEEE Trans. Commun.*, vol. COM-36, pp. 401-09, April 1988.

[24] M. V. Eyuboglu and S. U. Qureshi, "Reduced-State Sequence Estimation with Set Partitioning and Decision-Feedback," *IEEE Trans. Commun.*, vol. COM-36, pp. 13-20, Jan. 1988.

[25] A. Duel-Hallen and C. Heegard, "Delayed Decision-Feedback Equalization," *IEEE Trans. Commun.*, vol. COM-37, pp. 428-36, May 1989.

[26] P. Chevillat and E. Eleftheriou, "Decoding of Trellis-Encoded Signals in the Presence of Intersymbol Interference and Noise," *IEEE Trans. Commun.*, vol. COM-37, pp. 669-76, July 1989.

[27] M. V. Eyuboglu and S. U. Qureshi, "Reduced-State Sequence Estimation for Coded Modulation on Intersymbol Interference Channels," *IEEE J. Select. Areas. Commun.*, vol. SAC-7, pp. 989-95, August 1989.

[28] 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. SAC-7, pp. 877-92, Aug. 1989.

[29] A. R. Calderbank and L. H. Ozarow, "Non-Equiprobable Signalling on the Gaussian Channel," *IEEE Trans. Inform Theory*, vol. IT-36, pp. 726-40, July 1990.

[30] G. D. Forney, Jr., "Trellis Shaping," *IEEE Trans. Inform. Theory*, to appear Mar. 1992.

[31] M. V. Eyuboglu and G. D. Forney, Jr., "Trellis Precoding: Combined Coding, Precoding and Shaping for Intersymbol Interference Channels," *IEEE Trans. Inform. Theory*, to appear Mar. 1992.

[32] Motorola Information Systems, "A Single-Carrier Modem for V.fast: Description and Test Results," Contribution to CCITT SG XVII Special Rapporteur Group for Greater Than 14.4 Kbit/s Duplex Data Transmission over the PSTN, February 1991.

[33] R. Price, "Nonlinearly Feedback Equalized PAM Versus Capacity for Noisy Filter Channels," *Proc. ICC '72*, June 1972.

[34] R. deBuda, "Some Optimal Codes Have Structure," *IEEE J. Select. Areas Commun.*, vol. SAC-7, 893-99, August 1989.

Biography

G. David Forney, Jr. [F] received the B.S.E. degree in electrical engineering from Princeton University, Princeton, New Jersey, in 1961, and the M.S. and Sc.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge in 1963 and 1965, respectively. In 1965, he joined the Codex Corporation, and became a vice president and director in 1970. Dr. Forney was editor of the *IEEE Transactions on Information Theory* from 1970 to 1973. From 1982 to 1986, he was vice president and director of technology and planning of the Motorola Information Systems Group, Mansfield, Massachusetts. He currently is a vice president of the technical staff of Motorola, Inc., based at Codex.

M. Vedat Eyuboğlu (S'80-M'84-SM'91) received the B.S. degree in electrical engineering from Boğazici University, Istanbul, Turkey, in 1978, and M.S. and Ph.D. degrees in electrical engineering from Rensselaer Polytechnic Institute, Troy, New York, in 1980 and 1984, respectively. In November 1983, he joined Codex Corporation, a subsidiary of Motorola Inc., in Mansfield, Massachusetts, where he currently is manager of Transmission Research and Advanced Development. Dr. Eyuboğlu is a Science Advisory Board Associate (SABA) and Noble Fellow of Motorola and an adjunct professor at Northeastern University, Boston, Massachusetts.