

# Trellis and Convolutional Precoding for Transmitter-Based Interference Presubtraction

Wei Yu, *Member, IEEE*, David P. Varodayan, and John M. Cioffi, *Fellow, IEEE*

**Abstract**—This paper studies the combination of practical trellis and convolution codes with Tomlinson–Harashima precoding (THP) for the presubtraction of multiuser interference that is known at the transmitter but not known at the receiver. It is well known that a straightforward application of THP suffers power, modulo, and shaping losses. This paper proposes generalizations of THP that recover some of these losses. At a high signal-to-noise ratio (SNR), the precoding loss is dominated by the shaping loss, which is about 1.53 dB. To recover shaping loss, a trellis-shaping technique is developed that takes into account the knowledge of a noncausal interfering sequence, rather than just the instantaneous interference. At rates of 2 and 3 bits per transmission, trellis shaping is shown to be able to recover almost all of the 1.53-dB shaping loss. At a low SNR, the precoding loss is dominated by power and modulo losses, which can be as large as 3–4 dB. To recover these losses, a technique that incorporates partial interference presubtraction (PIP) within convolutional decoding is developed. At rates of 0.5 and 0.25 bits per transmission, PIP is able to recover 1–1.5 dB of the power loss. For intermediate SNR channels, a combination of the two schemes is shown to recover both power and shaping losses.

**Index Terms**—Broadcast channel, channels with noncausal side information, convolutional codes, dirty-paper precoding, shaping codes, Tomlinson–Harashima precoding (THP), trellis codes, trellis precoding (TP), trellis shaping (TS), vectored digital subscriber line (DSL).

## I. INTRODUCTION

CONSIDER a multiuser communications scenario in which a centralized transmitter wishes to transmit independent information to several remote users in geographically different locations at the same time. This downlink-transmission environment is often modeled as a broadcast channel. Because of the mutual interference caused by multiple signals, the capacity region and the optimal transmission strategy for the broadcast

channel have been long-standing open problems in multiuser information theory. Recently, however, a transmitter-based interference presubtraction scheme has gained much attention. The basic idea of presubtraction is as follows. Suppose that the transmitter wishes to send a signal  $x_1$  to the first user, and a signal  $x_2$  to the second user. Since both signals are completely known prior to transmission, the transmitter may presubtract  $x_2$  from  $x_1$ , so that the first receiver receives  $x_1$  as if the interference does not exist. Moreover, due to a surprising result known as “writing on dirty paper” [1], such a presubtraction scheme may be implemented without an increase in transmit power. Multiuser interference (MUI) precoding schemes were first proposed in [2] and [3]. Later, several independent results [4]–[6] further showed that MUI presubtraction in fact achieves the sum capacity of the multiantenna broadcast channel. Most recently, [7] fully established that the same is true for the entire capacity region.

The key theoretical insight for MUI presubtraction is the “writing on dirty paper” result due to Costa [1]. Consider a Gaussian noise channel corrupted by an additive interference signal  $S$  that is known to the transmitter but not to the receiver

$$Y = X + S + Z \quad (1)$$

where  $X$  and  $Y$  are the transmitted and received signals, respectively,  $S$  is the known interference, and  $Z$  is the unknown Gaussian noise. In a classical paper, Costa [1] showed that the capacity of this channel under a transmit power constraint is the same as if  $S$  does not exist, *provided that  $S$  is known noncausally at the transmitter*. For Costa’s result to hold, the transmitter has to know not only the present and past history of  $S$ , but also the future values of  $S$ . This setup, in fact, precisely models a multiuser downlink situation, where the interfering signal is the transmit signal to other users. Interference is known at the transmitter noncausally because the source information bits are typically stored in a buffer, and future values of the interference can be preconstructed from the buffered bits. The achievability proof of Costa’s result involves a coding technique known as binning. Costa’s paper is titled “Writing on dirty paper” because it models a transmitter which attempts to encode information on a piece of paper partially corrupted by dirt that is seen at the transmitter but not known at the receiver. Precoding methods for the dirty-paper channel are sometimes referred to as “dirty-paper precoding.”

The main practical motivation for this paper is a vectored digital subscriber line (DSL) application where crosstalk interference can be presubtracted at the transmitter. DSL is a family of high-speed data-transmission standards for the conventional telephone twisted pairs. Telephone wires are usually bundled

Paper approved by C. Schlegel, the Editor for Coding Theory and Techniques of the IEEE Communications Society. Manuscript received April 30, 2004. This work was supported in part by a Stanford Graduate Fellowship, in part by the Natural Science and Engineering Research Council (NSERC), in part by the Canada Research Chair Program, in part by Bell Canada University Laboratory, and in part by Communication and Information Technology Ontario (CITO). This paper was presented in part at IEEE Globecom, San Antonio, TX, December 2001.

W. Yu is with the Electrical and Computer Engineering Department, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: weiyu@comm.utoronto.ca).

D. P. Varodayan was with the Electrical and Computer Engineering Department, University of Toronto, Toronto, ON M5S 3G4, Canada. He is now with the Electrical Engineering Department, Stanford University, Stanford, CA 94305 USA (e-mail: varodayan@stanford.edu).

J. M. Cioffi is with the Electrical Engineering Department, Stanford University, Stanford, CA 94305 USA (e-mail: cioffi@stanford.edu).

Digital Object Identifier 10.1109/TCOMM.2005.851605

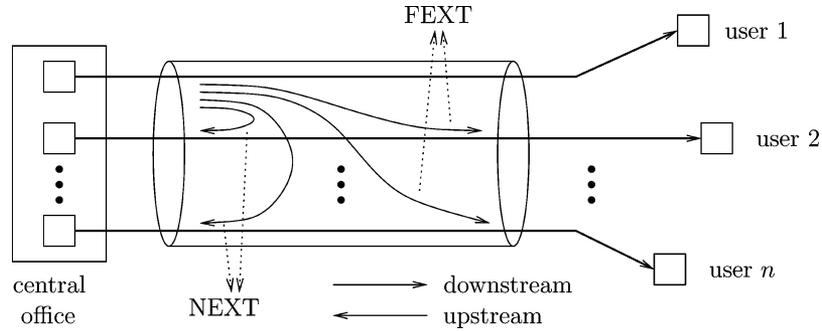


Fig. 1. Vectored DSLs.

together in a group of 50–100 lines at the central office. Because of their physical proximity, the crosstalk caused by the electromagnetic interference between the neighboring lines is often the dominating noise source in the transmission environment. Fig. 1 illustrates the nature of the crosstalk. The near-end crosstalk (NEXT) refers to crosstalk emitted into the receivers located at the same side as the transmitter. The far-end crosstalk (FEXT) refers to crosstalk emitted into the receivers located at the opposite end. Both NEXT and FEXT can have detrimental effects on DSL transmission.

If each DSL line transmits and receives information independently, the transmission environment must be modeled as an interference channel, where interference is typically regarded as noise. However, if multiple transmitters at the central office coordinate, then array signal-processing techniques may be used, and crosstalk cancelation becomes feasible. Because both NEXT and FEXT are completely known at the transmitter, the dirty-paper result implies that the capacity of the DSL channels can be as high as if the interference does not exist. This idea of interference presubtraction was first explored in [2], where a “vectored DSL” system is proposed. Vectored DSL uses Tomlinson–Harashima precoding (THP) to presubtract crosstalk, and it has the potential to double the transmission rate for short-range DSL services.

This paper studies practical coding methods for interference presubtraction at the transmitter. We use THP as a starting point. THP [8], [9] is a pre-equalization technique originally proposed for channels with intersymbol interference (ISI). It uses a modulo operation to presubtract interference with minimal power increase. However, compared with a channel with no interference, THP suffers from several sources of losses. As was pointed out in [10], precoding for the ISI channel incurs shaping loss in the high-signal-to-noise ratio (SNR) regime, and power and modulo losses in the low-SNR regime. The objective of this paper is to characterize the nature of similar losses in THP for MUI subtraction, and to propose practical ways to recover some of these losses for trellis- and convolutional-coded systems. As the main result of this paper shows, when combined with THP, convolutional and trellis codes designed for the usual additive white Gaussian noise (AWGN) channel can be adopted for the dirty-paper channel, with only a small performance loss and with a moderate amount of additional complexity.

The dirty-paper channel has been the subject of many studies in the information theory literature. In [11], Erez *et al.* consid-

ered the dirty-paper channel with causal side information and first pointed out the connection between THP and the dirty-paper result. They showed that at a high SNR, shaping loss is the only significant precoding loss, and at a low SNR, a partial interference-cancellation scheme can be implemented to recover some of the power and modulo losses. The focus of this paper is on practical coding schemes for the dirty-paper channel. One of the main contributions of this paper is to illustrate that trellis precoding (TP), previously developed for achieving a shaping gain in the ISI channel [12], [13], can be extended to the dirty-paper channel to recover most of the shaping loss in the high-SNR regime. Empirically, trellis shaping (TS) is demonstrated to be within 0.2 dB of the performance of the equivalent code on the AWGN channel.

In addition, this paper also investigates precoding methods for the low-SNR channel. It is previously known that in the low-SNR regime, partial interference presubtraction (PIP) outperforms total interference presubtraction. This paper studies practical implementation of PIP, and proposes a decoding metric that combines PIP with convolutional codes to recover some of the power and modulo losses in the low-SNR regime. The performance of the PIP is shown to be within 2.5 dB of the equivalent codes over the AWGN channel at rates of 0.25 and 0.5 bits per dimension. These simulation results are consistent with similar results in the low-density parity-check (LDPC) context [14]. Finally, we also present a scheme for the intermediate-SNR regime that combines shaping and PIP, and achieves within 0.5 dB as it would over the AWGN channel, thereby confirming Costa’s capacity result practically in that regime.

To completely recover all of the power, modulo, and shaping losses, a nested-coding approach can be used to implement the binning strategy. In this regard, Zamir *et al.* [15] showed that an inflated-lattice strategy can, in theory, achieve Costa’s capacity with a nested-lattice coding scheme. For practical coding systems, Erez and ten Brink [16] showed that nested repeat-accumulate codes and convolutional codes are effective. In general, however, true capacity-achieving practical coding for the dirty-paper channel is still an area of future research.

The rest of the paper is organized as follows. Section II explains the connection between THP and the dirty-paper channel, and illustrates the source of power, modulo, and shaping losses. Section III presents the trellis-shaping method, and illustrates its performance on high-SNR channels. Section IV presents the PIP scheme, and illustrates its performance on the low- and intermediate-SNR regimes. Section V contains concluding remarks.

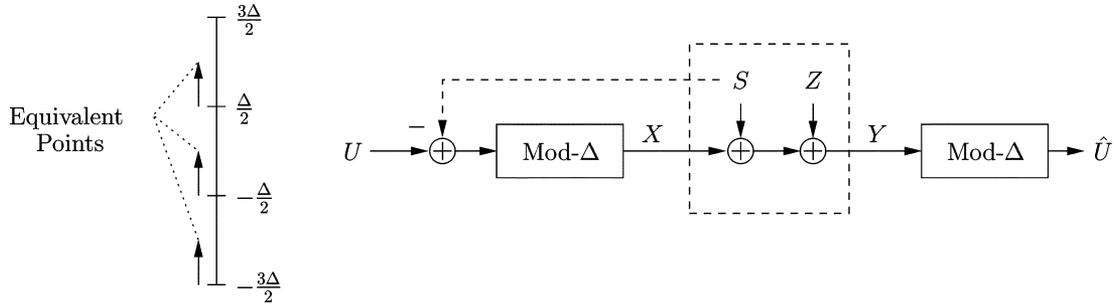


Fig. 2. THP for the dirty-paper channel.

## II. TOMLINSON–HARASHIMA PRECODING

### A. Modulo Precoding

It is well known that the capacity of an ordinary AWGN channel  $Y = X + Z$ , with an input power constraint  $P_X$  and a noise power  $P_Z$ , is

$$C = \frac{1}{2} \log_2(1 + \text{SNR}) \text{ bits/dim} \quad (2)$$

where SNR is defined as  $P_X/P_Z$ . Costa's result shows that the capacity does not change when the AWGN channel is corrupted by an interference signal with power  $P_S$ , if the interference is known noncausally at the transmitter. The realization of Costa's dirty-paper capacity involves a binning strategy, which may be implemented using nested codes [15]. However, at a high SNR, dirty-paper precoding may be approximated by THP.

THP [8], [9] was originally designed for the ISI channel, and it can be readily extended for the presubtraction of MUI at the transmitter. The main idea is shown in Fig. 2. The intended message is denoted  $U$ , and the interference signal is denoted  $S$ . The goal is to convey  $U$  to the receiver in the presence of  $S$ . Without loss of generality, assume that  $P_S$  is large. (When  $P_S$  is small, pseudorandom dithering can always be added to both the transmitter and the receiver to make the effective  $P_S$  large.) Since  $S$  is known at the transmitter at every time instant, in order to convey an intended symbol  $U$ , a precoder may send  $U' = U - S$  to compensate for the interference. However, when  $S$  is large, the power of  $U'$  may be exceedingly large. The main idea of THP is to constrain the intended symbol  $U$  to be within a finite interval  $[-\Delta/2, \Delta/2)$ . Instead of sending  $U'$ , the encoder sends  $U'$  modulo- $\Delta$ . In effect, all  $U'$ 's that differ by an integer multiple of  $\Delta$  are regarded as the same symbol. (See Fig. 2.) In addition, the decoder recognizes the same equivalence relation by implementing a modulo- $\Delta$  operation. In the absence of noise, this equivalence relation allows  $U$  be completely recovered at the receiver.

Essentially, the modulo- $\Delta$  operation constrains the transmitted signal  $X$  to be within the finite interval  $[\Delta/2, \Delta/2)$ . When  $P_S$  is large,  $X$  is approximately uniformly distributed in the interval. In this case, the transmit power is reduced to approximately  $\Delta^2/12$ , which is independent of the interference power. This is all accomplished with the distinguishability  $U'$  maintained.

The encoding process in THP can also be thought of as the process of modifying  $S$  to one of the equivalent representatives of the symbol. This is a one-dimensional (1-D) quantization

process. The modulo- $\Delta$  geometry ensures that there is one representative within distance  $\Delta/2$  from every possible  $S$ . As will be seen later, the THP can be interpreted as a 1-D implementation of "writing on dirty paper." One of the main contributions of this paper is a practical trellis-precoding scheme that generalizes THP to high dimensions.

### B. Precoding Losses

When compared with the performance of the standard pulse-amplitude modulation (PAM) on an AWGN channel, THP incurs some nonnegligible performance losses. The precoding loss was first pointed out by Shamai and Laroia [10] for the ISI channel. They characterized three sources of loss as shaping loss, power loss, and modulo loss. This characterization is equally applicable to the MUI presubtraction scenario being considered in this paper.

- *Shaping Loss:* First of all, because of the use of the modulo operation, the input signal  $U$  must be amplitude-limited. In particular, when  $S$  is large,  $X$  is uniformly distributed in  $[-\Delta/2, \Delta/2)$ . Such a uniform shape introduces the so-called shaping loss. Shaping loss is due to the fact that Shannon's capacity formula for AWGN channels requires the capacity-achieving input distribution to be a Gaussian distribution. Signaling using a uniform distribution introduces the finite shaping loss. The shaping loss due to  $M$ -ary PAM ( $M$ -PAM) over AWGN channels is insignificant at low SNR, but tends to the ultimate shaping loss of 1.53 dB for high SNR. In fact, as shown by Erez *et al.* [11], shaping loss is intimately related to the fact that THP only takes the present value of the interference signal  $S$  into account, and does not take advantage of the knowledge of future interference. This observation is the motivation for the main result of this paper, which is a trellis-shaping code that is capable of recovering shaping loss based on the entire interfering sequence.
- *Power Loss:* Second, the transmitted signal  $X$  in THP can have more power than the intended signal  $U$ . This difference is referred to as the power loss. Presubtracting Gaussian-distributed  $S$  from  $U$  and applying a modulo operation spreads  $X$  continuously over the modulo region. Assume that  $U$  is an equiprobable  $M$ -PAM with constellation points in  $\{((-M + 1)/2M)\Delta, ((-M + 3)/2M)\Delta, \dots, ((M - 1)/2M)\Delta\}$ . When  $S$  has a sufficiently large

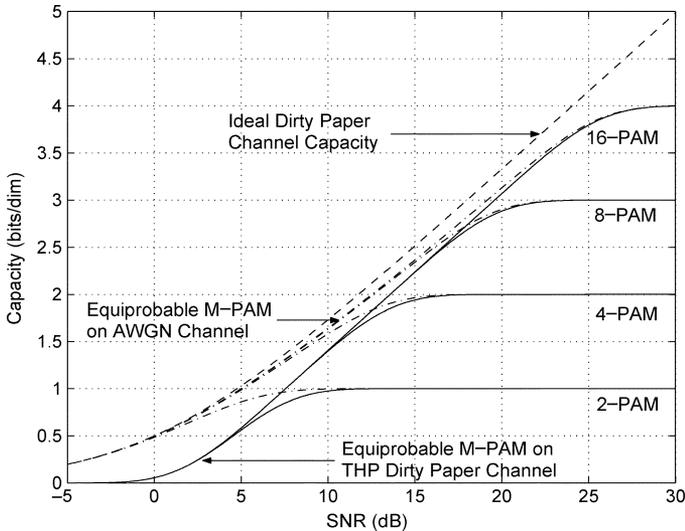


Fig. 3. Capacity of ideal dirty-paper channel compared with  $M$ -PAM capacities of AWGN and THP dirty-paper channels.

power,  $X$  would also be uniform over the modulo region  $[-\Delta/2, \Delta/2)$ . Such an  $X$  has more energy than  $U$ , so THP introduces a power loss at the transmitter equal to  $P_X/P_U$ . A simple calculation shows that for the  $M$ -PAM constellation, the power loss equals  $M^2/(M^2 - 1)$ . Power loss is significant only for low-SNR channels, where small constellations are used.

- **Modulo Loss:** The modulo operation at the THP receiver introduces a different loss, called modulo loss. Observe that the receiver collapses all received symbols that differ by  $\Delta$  to the same value, so they can be viewed as multiple representations of the same symbol. With THP, received symbols at the boundary of a constellation may now be mistaken for symbols at the opposite boundary of the constellation, thus incurring potential errors. Modulo loss is related to the proportion of constellation points at the boundaries, and is more pronounced for small constellations (on low-SNR channels). Notice also that channel-coding schemes for the THP dirty-paper channel must take this wrapping effect into account. For convolutional- and trellis-coded systems, this alters the calculation of branch costs in the Viterbi algorithm, which increases the number of nearest-neighbor error patterns.

Fig. 3 compares the capacity of the dirty-paper channel given in (2) with the equiprobable  $M$ -PAM capacities of the AWGN and THP dirty-paper channels. The capacity of the THP dirty-paper channel is computed as in [17]

$$I(\hat{U}; U) = h(\hat{U}) - h(\hat{U}|U) \quad (3)$$

$$= h(Y \bmod \Delta) - h(Z \bmod \Delta). \quad (4)$$

Notice that the performance loss of an  $M$ -PAM-coded system on the AWGN channel is purely shaping loss, while the THP dirty-paper channel suffers from additional power and modulo losses, which can be large in the low-SNR regime. The goal of the rest of the paper is to present various approaches to recover these precoding losses in various SNR regimes.

### III. TRELLIS CODING FOR THE DIRTY-PAPER CHANNEL

#### A. Trellis Precoding

At high SNR, THP on the dirty-paper channel incurs a negligible loss, compared with an  $M$ -PAM constellation on the AWGN channel. However, it suffers a shaping loss, compared with the ideal dirty-paper capacity. The main objective of this section is to present a trellis-precoding scheme that is capable of recovering most of the shaping loss in the high-SNR regime. The key insight, first pointed out in [11], is that the THP presubtracts only the current  $S$ , and does not consider future values of  $S$ . With only causal side information, the modulo operation has to be done on a symbol-by-symbol basis, producing an output that is uniformly distributed between  $-\Delta/2$  and  $\Delta/2$ . This uniform distribution corresponds to a high-dimensional cubic shape, thus incurring a shaping loss when compared with the spherical shape of an optimal Gaussian code. To recover the shaping loss, noncausal side information must be used to perform the modulo operation on high-dimensional spheres. The modulo operation in THP can be thought of as a 1-D quantization process, where the output is the quantization error. The generalization of the high-dimensional modulo operation is, therefore, a vector quantizer which outputs the vector quantization error.

A conceptual model for a vector quantization-based precoder is shown in Fig. 4. It works as follows. First, a codeword sequence  $U^n$  is generated by the encoder of an error-correcting code. The additive interference sequence  $S^n$  is presubtracted from the codeword, and the difference sequence is then quantized by a vector quantizer. The quantization error is sent as the input to the channel. The channel adds the interference and noise. At the decoder, the received sequence is first quantized by the same vector quantizer. The quantization error is sent to the decoder of the error-correcting code to recover the message. Fig. 5 illustrates the relation between  $U^n$ ,  $S^n$ , and  $X^n$ . The spheres denote the Voronoi regions associated with the quantizer outputs. The codeword sequence  $U^n$  is designed to be confined within the Voronoi region. Each codeword is given multiple equivalent representatives corresponding to multiple quantizer outputs. The equivalent representatives are illustrated by the arrows in Fig. 5. Since the interference sequence  $S^n$  is known noncausally, the precoder can construct an input sequence  $X^n$  to steer  $S^n$  to the closest representative of  $U^n$ . It is easy to see that as long as the codeword is confined to the Voronoi region, perfect reconstruction is possible in the absence of noise. This vector-quantization approach can be viewed as a generalization of THP. The 1-D modulo- $\Delta$  operation is replaced by a vector quantizer which performs a modulo operation with respect to a Voronoi region. After the modulo operation, the precoder outputs are approximately uniformly distributed in the Voronoi region. The Voronoi region of an ideal vector quantizer is a high-dimensional sphere, thus achieving a shaping gain. In other words, the transmit power may be reduced by up to 1.53 dB when compared with a cubic shape.

The idea of applying Voronoi constellations to lattice coding and quantization is due to Conway and Sloane [18]. Forney [19] proposed the use of the Voronoi region of a lattice to achieve a

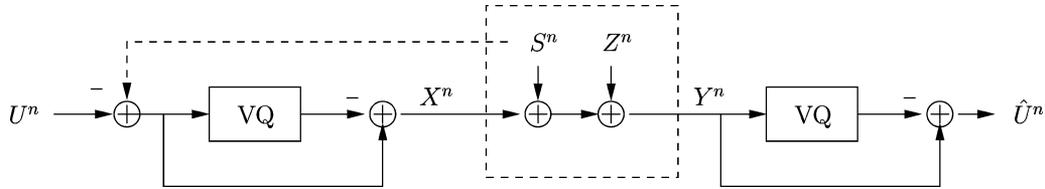


Fig. 4. Precoding via vector quantization.

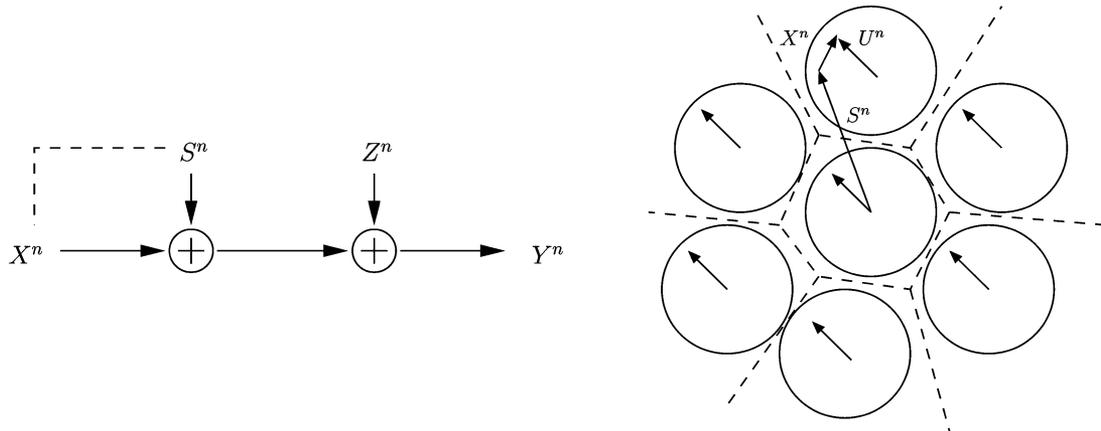


Fig. 5. Precoding with respect to a Voronoi region.

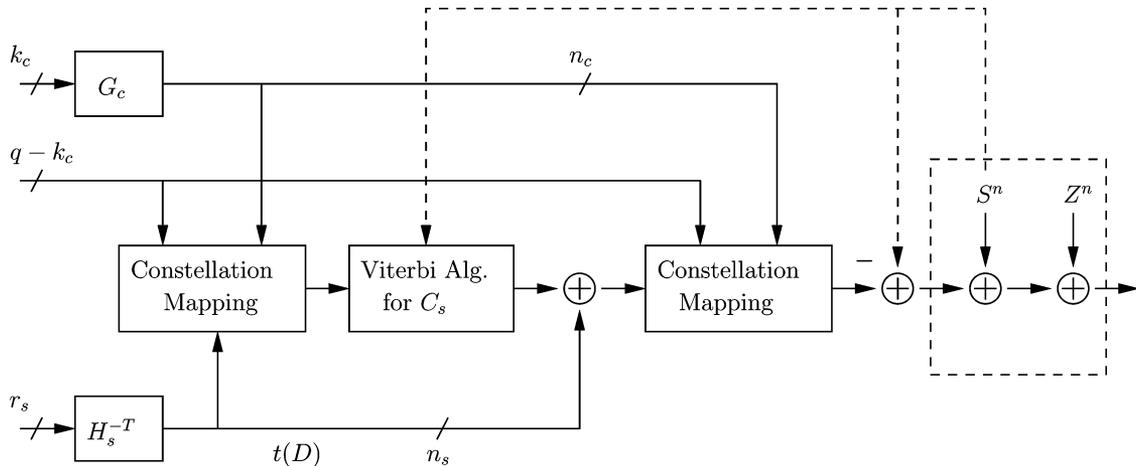


Fig. 6. Implementation of Voronoi precoder.

shaping gain on the AWGN channel. This idea was further extended in [12], where a trellis-coded quantizer [20] was used to achieve shaping gain. In an AWGN channel, the capacity-achieving distribution is a Gaussian distribution, which is equivalent to a uniform distribution over an  $n$ -dimensional sphere over a block length of  $n$ . However, traditional rectangular constellation-based trellis codes cover the  $n$ -cube uniformly, thus suffering from a shaping loss up to 1.53 dB. To recover the shaping loss, Forney [12] proposed first to expand the constellation size slightly, then to modulo the expanded sequence with respect to the Voronoi region of a trellis-shaping code. It is clear that shaping for the AWGN channel is very similar to shaping for a channel with side information. In fact, Eyuboglu and Forney [13] further extended TS by combining shaping with 1-D THP and trellis coding for the ISI channel.

We now describe a trellis-precoding scheme for channels with side information. The operation of the precoder is shown in Fig. 6. The following system description relies on the interpre-

tation of a trellis code as a collection of paths through the constellation. These paths have a good minimum distance property, and they can be represented by a finite state machine, thus allowing efficient decoding. Such a collection of trellis paths can be used as reconstruction values in a vector quantizer.

Two codes are working independently.  $C_c$  is a trellis-channel code, with  $G_c$  as its  $k_c/n_c$  convolutional coset encoder.  $C_s$  is a  $k_s/n_s$  trellis-shaping code, whose Voronoi region will be used as the basis of a modulo operation. The input bits are divided into three groups. At any given time instance, the trellis-channel code accepts  $q$  bits, of which  $k_c$  bits are encoded into coset select bits by the convolutional code  $G_c$ . The other  $q - k_c$  signal select bits remain uncoded, and select constellation points within each coset. The output of the trellis-channel code is over a constellation of size  $2^{q-k_c+n_c}$ .

In the shaping code, the output constellation of the trellis-channel code is repeated  $2^{n_s}$  times in two dimensions, resulting in  $2^{n_s}$  nonoverlapping regions. A sequence of these regions is

selected by a third group of  $r_s = n_s - k_s$  bits entering an inverse syndrome former  $H_s^{-T}$  for the shaping code  $C_s$ . The  $n_s$ -bit output of  $H_s^{-T}$  selects one of the  $2^{n_s}$  regions. Over time, the inverse syndrome former produces a sequence  $t(D)$  of  $n_s$ -bit symbols. Via constellation mapping, the overall codeword can be thought of as the combination of  $2^q$  trellis paths in each constellation region, with the sequence of regions  $t(D)$ .

Now, the modulo-like role of the shaping code operates as follows. The information bits are contained in the  $r_s$ -bit stream. Further, all sequences of regions with the same syndrome under  $C_s$  are deemed equivalent. (In other words, two sequences of regions are equivalent if their difference is a codeword of  $C_s$ .) So, the shaping encoder is free to add any codeword of  $C_s$  to  $t(D)$ , while preserving the  $r_s$  information bits. The choice of this codeword is made by the Viterbi algorithm for  $C_s$ . The criterion for such a selection can, in principle, be anything as far as decoding is concerned. To minimize transmit power for channels with side information, the criterion here is the minimization of the square difference between the overall codeword and the entire side-information sequence  $S^n$ . To summarize, the Viterbi algorithm for  $C_s$  compares  $S^n$  with all possible overall codewords determined jointly by the trellis encoder for  $C_c$  and the syndrome sequence for  $C_s$ , and outputs a codeword of  $C_s$ . The desired overall codeword is now obtained in two steps. This codeword of  $C_s$  is added to  $t(D)$  to form a new sequence of regions, which is then combined with the codeword of  $C_c$  through constellation mapping. Finally, the side information  $S^n$  is pre-subtracted, and it is the low-power quantization error  $X^n$  that is transmitted. On the decoder side, the bits are recovered with a usual trellis decoder for  $C_c$  and a syndrome mapper for  $C_s$ . The decoder is identical to the one in [13].

There are two key differences between this trellis precoder for a channel with side information and a trellis-shaping precoder for an ISI channel, as described in [13]. In [13], a shaping code selects the region sequence so that the output codeword has the minimum energy, while here, a precoder selects the region sequence so that it takes the minimum energy to steer the side information to a correct codeword. The second difference is more subtle. In a shaping code for an ISI channel, a small amount of constellation expansion suffices, thus a rate-1/2 shaping code is sufficient. For precoding for a channel with side information, it is desirable to have a shaping code rate such as 3/4 or 5/6 to force greater constellation expansion. The reason for expansion is to ensure that the side-information sequence lies entirely within the expanded constellation. In practice, when the magnitude of the side-information sequence is not known in advance, it is necessary to add a modulo- $M$  operation outside of the expanded constellation. If the expansion is insufficient with respect to  $S$ , the output distribution becomes uniform inside a cubic shape, and shaping loss may not be recovered. Note that the actual transmit symbols have much smaller amplitudes, compared with the expanded constellation. Since only the difference between the intended signal and the side-information sequence is transmitted, the expanded constellation does not affect the dynamic range requirement of a practical transmitter.

The shaping gain for the precoder depends on the shape of the Voronoi region of the shaping code. Therefore, the trellis-shaping codes in [12] can be used directly for multiuser

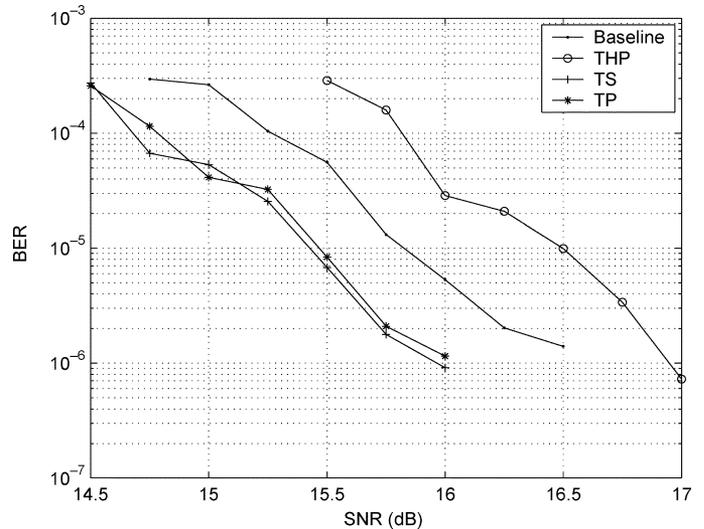


Fig. 7. BER versus SNR at 2 bits/dim using 64-state channel codes and 8-state shaping codes.

precoding with exactly the same shaping gain. In particular, a simple four-state trellis-shaping code already achieves almost 1 dB shaping gain.

As far as the transmit power is concerned, trellis shaping is capable of recovering almost all of a 1.53 dB shaping loss [12]. However, as noted by Fischer *et al.* [21], [22], a trellis-shaping receiver can suffer from error propagation, since errors in decoding the channel-coded stream affect the uncoded shaping stream as well. The alternative method suggested in [21] selects the initial shaping sequence using an ISI precoder. However, as the dirty-paper channel has no ISI, this remedy is not directly applicable here. Fortunately, as seen in the simulation results, the effect of error propagation is not significant for the dirty-paper channel. Finally, we note that the recent work of Fischer *et al.* [23] took a related approach that used a lattice-shaping coding instead of a trellis-shaping code.

## B. High-SNR Simulation Results

The performance of trellis precoding is evaluated on the dirty-paper channel. Simulation results for four different schemes are presented for the purpose of comparison: the baseline trellis-coded modulation on AWGN channels; the TS scheme over the AWGN channel; the THP scheme over the dirty-paper channel; and the TP scheme with TS over the dirty-paper channel. Figs. 7 and 8 depict the performance of implementations operating at 2 and 3 bits/dim, respectively.

In Fig. 7, the TS and TP schemes use a 16-quadrature amplitude modulation (QAM) constellation, divided into eight cosets with two constellation points in each coset. The channel trellis code  $C_c$  consists of one uncoded signal select bit and three coded coset select bits encoded with a 64-state, rate-2/3 systematic convolutional encoder  $G_c$ . The generator matrix of  $G_c$  is

$$\begin{bmatrix} 1 & 0 & \frac{h_1(D)}{h_0(D)} \\ 0 & 1 & \frac{h_2(D)}{h_0(D)} \end{bmatrix} \quad (5)$$

with  $h_2(D) = D + D^2 + D^3 + D^4$ ,  $h_1(D) = D + D^3 + D^5$ , and  $h_0(D) = 1 + D^2 + D^3 + D^6$ . To provide shaping, the

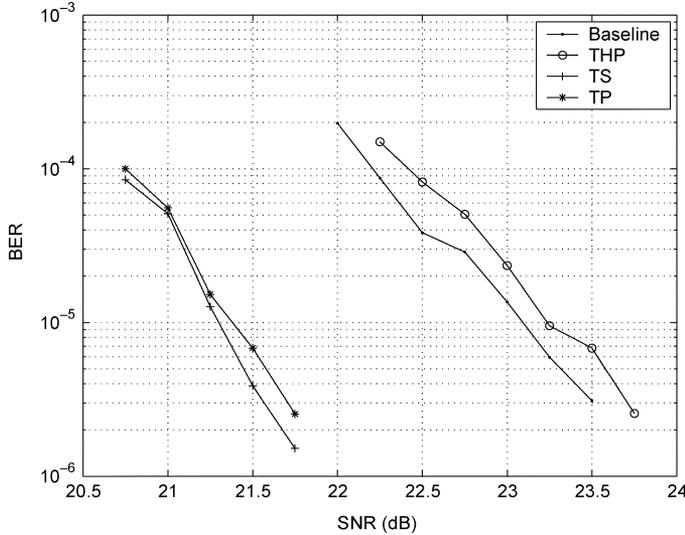


Fig. 8. BER versus SNR at 3 bits/dim using 64-state channel codes and 8-state shaping codes.

basic 16-QAM constellation is repeated 64 times. The 64 regions form an outer 64-QAM constellation, which is subdivided into eight cosets with eight constellation points in each coset. The shaping trellis code  $C_s$  consists of three uncoded signal select bits and three coded coset select bits encoded by an 8-state, rate-2/3 systematic convolutional code (with generator matrix as in (5) with  $h_2(D) = D^2$ ,  $h_1(D) = D$ , and  $h_0(D) = 1 + D^3$ ). If viewed as a channel code,  $C_s$  would produce six output bits per time instant, so  $n_s = 6$ . The complete Voronoi precoder parameters (see Fig. 6) are  $(q, k_c, n_c, r_s, n_s) = (3, 2, 3, 1, 6)$ . The corresponding baseline and THP systems use a 32-QAM constellation, which is divided into eight cosets with four constellation points in each coset. The channel trellis code consists of the same 64-state, rate-2/3 systematic convolutional encoder  $G_c$  as above, but with two uncoded signal select bits at a time.

In Fig. 8, the TS and TP schemes use a 64-QAM constellation divided into eight cosets with eight constellation points in each coset. The channel trellis code  $C_c$  consists of three uncoded signal select bits and three coded coset select bits encoded by a 64-state, rate-2/3 systematic convolutional encoder  $G_c$  (with generator matrix as in (5) with  $h_2(D) = D + D^2 + D^3 + D^4$ ,  $h_1(D) = D + D^3 + D^5$ , and  $h_0(D) = 1 + D^2 + D^3 + D^6$ ). Again, to provide shaping, the basic 64-QAM constellation is repeated 64 times. The 64 regions form an outer 64-QAM constellation, which is subdivided into eight cosets. The shaping trellis code  $C_s$  is the same as in the previous simulation. It consists of three uncoded signal select bits and three coded coset select bits encoded by an 8-state, rate-2/3 systematic convolutional encoder (with generator matrix as in (5) with  $h_2(D) = D^2$ ,  $h_1(D) = D$ , and  $h_0(D) = 1 + D^3$ ). The complete Voronoi precoder parameters (see Fig. 6) are  $(q, k_c, n_c, r_s, n_s) = (5, 2, 3, 1, 6)$ . The corresponding baseline and THP systems use a 128-QAM constellation, which is divided into eight cosets with 16 constellation points in each coset. The channel trellis code consists of the same 64-state, rate-2/3 systematic convolutional encoder  $G_c$  as above, but with four uncoded signal select bits.

Although the channel codes used in all four schemes are similar, they differ slightly in the number of uncoded signal select bits ( $q - k_c$ ), and hence, the size of the basic constellation ( $2^{q-k_c+n_c}$ ). Note also that TP over the dirty-paper channel has the same implementation complexity as the TS scheme over the AWGN channel. Each data point was obtained by simulation with blocks of  $10^5$  bits until at least 200 bit errors accrued, with SNR measured as the ratio of power of the simulated complex transmitted signal to power of the simulated complex noise signal.

The key observation is that the TP scheme on the dirty-paper channel performs almost as well as the usual TS scheme over the AWGN channel. The performance discrepancy of less than 0.2 dB is a small power loss, due to the fact that the transmitted signal for TS is discrete, but for TP, it is continuous. With respect to capacity, performance using 64-state channel codes is approximately 3.75 dB away at bit-error rate (BER) =  $10^{-5}$  for the transmission rate of 2 bits/dim; it is 3.25 dB away at the same BER for 3 bits/dim.

#### IV. CONVOLUTIONAL CODING FOR THE DIRTY-PAPER CHANNEL

##### A. Partial Interference Presubtraction

We have so far dealt with the high-SNR channel where shaping loss is the main concern. For the low-SNR dirty-paper channels, the power and/or modulo loss of the THP must also be taken into account. In this section, we use a PIP scheme as an effective way to partially recover the power loss at a low SNR. This scheme draws from Costa's information theoretical proof of the dirty-paper channel capacity, and is a 1-D implementation of Erez *et al.*'s reinterpretation in the context of the inflated lattice strategy [11]. The main contribution of this section is in showing how partial presubtraction may be incorporated in the Viterbi decoding of convolutional codes.

At a low SNR,  $U$  is a binary signal. The THP loss for a binary signal comes from two sources. First, precoding creates a uniformly distributed signal, which has a significantly higher power than a binary signal. In addition, compared with the usual uncoded 2-PAM signal on an AWGN channel, where each constellation point has exactly one nearest neighbor, with THP, the number of nearest neighbors increases to two. This is also significant and it leads to a modulo loss. The PIP scheme, shown in Fig. 9, is designed to recover some of these losses.

The main idea of partial interference subtraction is to presubtract  $\alpha S$  instead of  $S$ , where  $\alpha \in (0, 1]$ . At the other end, the receiver multiplies the received signal by  $\alpha$  before applying the modulo operation. Note that

$$\hat{U} = (\alpha X + \alpha S + \alpha Z) \bmod \Delta \quad (6)$$

$$= (U - (1 - \alpha)X + \alpha Z) \bmod \Delta \quad (7)$$

(since  $\alpha S = (U - X) \bmod \Delta$ ). The effective noise is thus the modulo of the weighted sum of two independent random variables, the uniform  $X$  and the Gaussian  $Z$ . By choosing  $\alpha < 1$ , the combined noise can be made smaller than the Gaussian noise  $Z$  alone. In fact, the optimal  $\alpha$  is the value that maximizes  $I(\hat{U}; U)$ . However, in practice, an  $\alpha$  that minimizes the power of  $(1 - \alpha)X + \alpha Z$  is also found to work well.

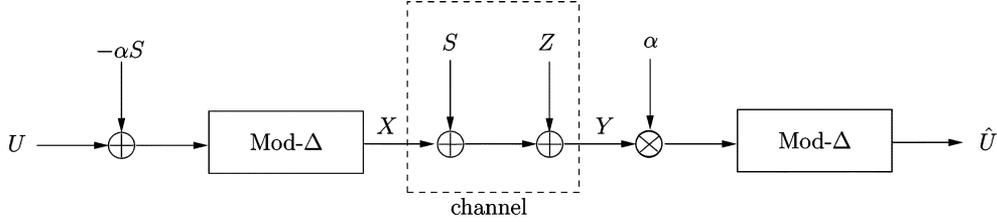
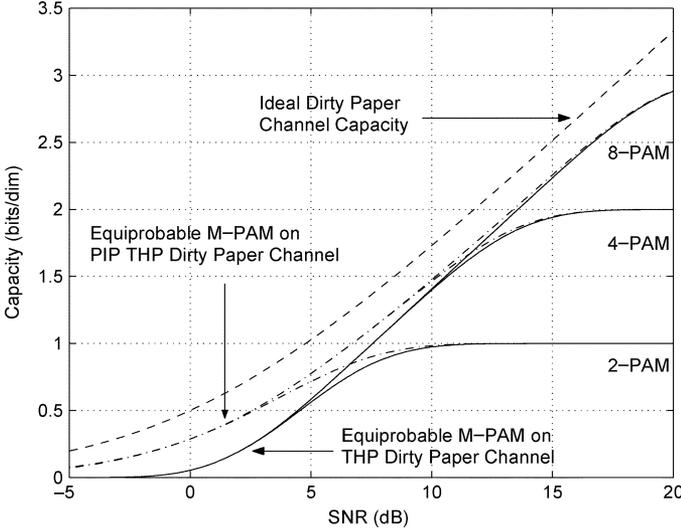


Fig. 9. Partial interference presubtraction in Tomlinson–Harishima precoding.


 Fig. 10. Capacity of ideal dirty-paper channel compared with equiprobable  $M$ -PAM capacities of THP and PIP THP dirty-paper channels.

The choice of  $\alpha$  that minimizes the total noise power turns out to coincide with Costa’s choice based on information theoretical considerations:  $\alpha = P_X/(P_X + P_Z)$  [11]. Fig. 10 quantifies the gain due to partial presubtraction, by comparing the capacity of the dirty-paper channel given in (2) with the equiprobable  $M$ -PAM capacities of the THP dirty-paper channel and the PIP THP dirty-paper channel. At 0.5 bits per transmission or below, the gain is over 2 dB, which is significant.

The main point of this section is to illustrate how PIP may be combined with convolutional codes. The decoding process is now slightly different when  $\alpha < 1$  is chosen. The usual Viterbi decoding of the convolutional code over the AWGN channel uses a Euclidean distance metric. However, under PIP, the effective noise is  $(1 - \alpha)X + \alpha Z$ . Since  $X$  is the output of a modulo operation, it is approximately independent and identically distributed (i.i.d.) with a uniform distribution. Thus, the effective noise is a combination of Gaussian and uniform components. Decreasing  $\alpha$  increases the power of the uniform part, and flattens the overall noise distribution. Therefore, at  $\alpha < 1$ , the usual Euclidean distance metric in the Viterbi algorithm gives too much preference to small-magnitude noise over large-magnitude noise. Instead, the following new metric should be used:

$$d(x, y) = -\log f(x - y) \quad (8)$$

where  $f(x)$  is the pdf of the noise  $(-(1 - \alpha)X + \alpha Z) \bmod \Delta$ . This metric captures exactly the self-information of the noise, and the Viterbi algorithm minimizes it over the whole received signal. In terms of complexity, the new metric is no worse than

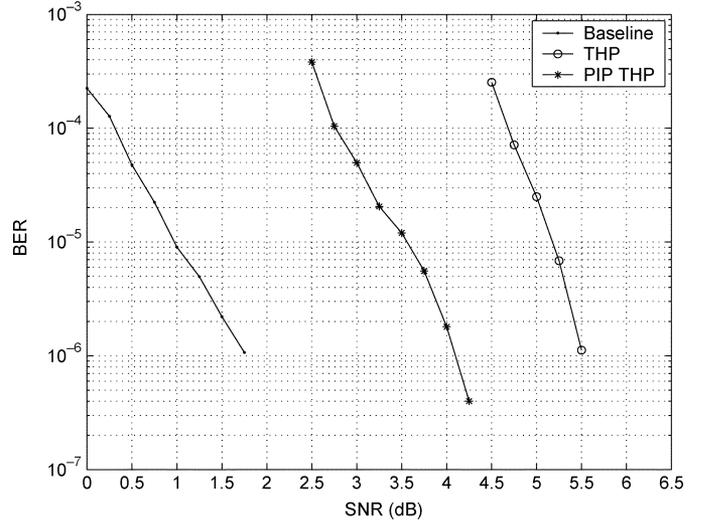


Fig. 11. BER versus SNR at 0.25 bits/dim, using 64-state convolutional codes.

the Euclidean one, since both are ultimately implemented as table lookups.

The need for the modified metric highlights the difficulty in truly achieving the capacity of the dirty-paper channel in the low-SNR regime using newer iterative decoding methods, such as the sum-product algorithm (see [24] and references therein). As SNR goes to zero, the optimal  $\alpha$  also goes to zero. This means that the noise in the channel is essentially uniform. Turbo and LDPC codes with iterative decoding are effective in the Gaussian noise channels, but the decoding problem for a uniform noise channel is much more difficult. Decoding with uniform noise is equivalent to quantization. Such a difficulty does not exist for Viterbi-based decoders.

### B. Low-SNR Simulation Results

Figs. 11 and 12 show the performance of baseline convolutional codes over AWGN channels, and the same codes with THP and PIP THP over dirty-paper channels, for transmission rates of 0.25 and 0.5 bits/dim, respectively.

The three schemes in Fig. 11 use identical 64-state, rate-1/4 convolutional codes (with 2-PAM), specified by generator matrix

$$[g_1(D) \quad g_2(D) \quad g_3(D) \quad g_4(D)] \quad (9)$$

where  $g_1(D) = 1 + D + D^4 + D^5 + D^6$ ,  $g_2(D) = 1 + D + D^2 + D^5 + D^6$ , and  $g_3(D) = g_4(D) = 1 + D^2 + D^3 + D^4 + D^6$ . The schemes in Fig. 12 also use identical 64-state, rate-1/2 convolutional codes (with 2-PAM), specified by generator matrix

$$[g_1(D) \quad g_2(D)] \quad (10)$$

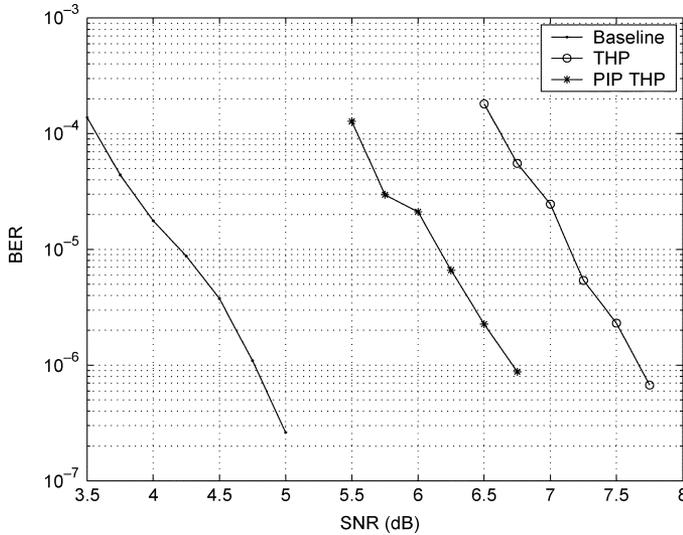


Fig. 12. BER versus SNR at 0.5 bits/dim, using 64-state convolutional codes.

where  $g_1(D) = 1 + D^2 + D^3 + D^5 + D^6$  and  $g_2(D) = 1 + D + D^2 + D^3 + D^6$ . Each data point was obtained by simulation with blocks of  $10^5$  bits until at least 200 bit errors accrued, with SNR measured as the ratio of power of the simulated complex transmitted signal to the power of the simulated complex noise signal.

The plots indicate that PIP recovers a sizeable portion of the total THP loss for low-SNR channels. The remaining 2–3 dB loss closely matches the modulo loss in Fig. 10. We also note that the performance gap to capacity of these PIP THP schemes is approximately 7 dB at  $\text{BER} = 10^{-5}$  for a transmission rate of 0.25 bits/dim, and 6 dB at the same BER for 0.5 bits/dim. In effect, about 2–2.5 dB of the gap to capacity is due to modulo loss, and 4 dB is due to the suboptimality of the 64-state channel codes. In other words, PIP is within 2.5 dB of an equivalent code on AWGN channels. This confirms Costa's capacity result for the low-SNR dirty-paper channel for practical codes.

### C. Intermediate-SNR Simulation Results

At intermediate SNR, both power loss and shaping loss are present, as shown in Fig. 3. To recover as much of the loss as possible, we combine TP and PIP. This scheme differs very slightly from the TP shown in Fig. 6 and described in the previous section. At the transmitter, the Viterbi algorithm for the shaping code  $C_s$  helps select an equivalent representative that lies close to  $\alpha S$ , from which  $\alpha S$  is presubtracted. At the receiver, the received signal is multiplied by  $\alpha$  before processing continues as per the decoder in [13]. Notice that with the use of a shaping code, the statistics of the transmitted signal is almost Gaussian, so the effective noise is also almost Gaussian. Therefore, the Euclidean metric is almost optimal, and it is not necessary to alter the cost metric in the Viterbi algorithm for the channel code.

Fig. 13 shows the performance of baseline trellis coding and TS over the AWGN channel, and THP and PIP TP over the dirty-paper channel, at a transmission rate of 1 bit/dim. The TS and PIP TP schemes use a 4-QAM channel trellis code  $C_c$  which has

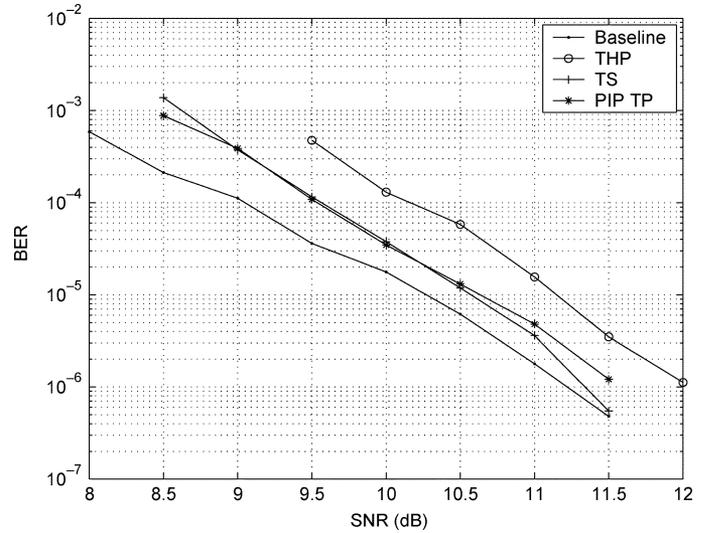


Fig. 13. BER versus SNR at 1 bit/dim, using 64-state channel codes and 8-state shaping codes.

a 64-state, rate-1/2 systematic convolutional encoder  $G_c$ , but no uncoded signal select bits. The generator matrix of  $G_c$  is

$$\begin{bmatrix} 1 & h_1(D) \\ & h_0(D) \end{bmatrix} \quad (11)$$

with  $h_1(D) = D^2 + D^3 + D^5$  and  $h_0(D) = 1 + D^4 + D^5 + D^6$ . For shaping, the basic 4-QAM constellation is repeated 64 times to form an outer 64-QAM constellation of regions, which is further subdivided into eight cosets with eight points in each coset. The shaping trellis code  $C_s$  consists of three uncoded signal select bits and three coded coset select bits encoded by an eight-state, rate-2/3 systematic convolutional encoder (with generator matrix as in (5) with  $h_2(D) = D^2$ ,  $h_1(D) = D$ , and  $h_0(D) = 1 + D^3$ ). The complete Voronoi precoder parameters (see Fig. 6) are  $(q, k_c, n_c, r_s, n_s) = (1, 1, 2, 1, 6)$ . The corresponding baseline and THP systems use an 8-QAM constellation divided into four cosets with two constellation points in each coset. The channel trellis code consists of the same 64-state, rate-1/2 systematic convolutional encoder  $G_c$  as above, but with one uncoded signal select bit. Each data point was obtained by simulation with blocks of  $10^5$  bits until at least 200 bit errors accrued, with SNR measured as the ratio of power of the simulated complex transmitted signal to the power of the simulated complex noise signal.

It can be seen from the figure that PIP TP performs within 0.5 dB of the baseline case. At  $\text{BER} = 10^{-5}$ , this 64-state trellis method has a gap to capacity of 5.75 dB. Using the theoretical plots in Fig. 10, we estimate that 2 dB of the gap is due to modulo loss, leaving 3.75 dB resulting from the suboptimality of the 64-state channel code.

## V. CONCLUSIONS AND DISCUSSIONS

This paper studies the combination of practical trellis and convolutional coding with Tomlinson–Harashima precoding (THP) for the presubtraction of multiuser interference (MUI) at the transmitter. THP (previously designed for the ISI channel) incurs significant precoding losses, compared with

the dirty-paper channel capacity. The precoding loss can be characterized as shaping loss, power loss, and modulo loss. Shaping loss dominates in the high-SNR regime. Power and modulo losses dominate in the low-SNR regime.

The main contribution of this paper is a trellis precoding (TP) method that is capable of almost completely recovering shaping loss for high-SNR channels with known interference. TP is based on the combination of THP and TS. The modulo operation in THP is generalized to a vector-quantization process. By presubtracting the future sample path of the interfering signal (rather than the symbol-by-symbol subtraction), the transmit signal can be shaped to approximate a Gaussian distribution, thus recovering the shaping loss. Simulation results show that at rates 1, 2, and 3 bits/dim, TP is within 0.5 dB of the performance of an equivalent code for AWGN channels.

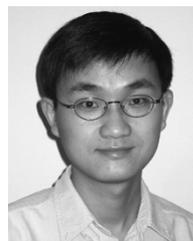
This paper also studies the low-SNR dirty-paper channels, and proposes a technique to combine partial interference presubtraction (PIP) with convolutional coding to recover the power loss. The main idea is to subtract interference partially, and to code for an effective noise that is a combination of self-noise and Gaussian noise, but has a smaller overall variance. The implementation of PIP requires a modification of the Viterbi decoding in convolutional codes. Simulation results show that it is within 2.5 dB of the performance of an equivalent code for AWGN channels at rates 0.25 and 0.5 bits/dim.

The basis of comparison in this paper is between equivalent codes on channels with or without known interference. Although more advanced coding strategies, such as turbo codes and LDPC codes, are expected to outperform the trellis codes and convolutional codes presented here, the comparison concerning the effectiveness of TP and PIP is expected to be valid, even when more advanced error-correcting codes are used. In particular, we expect the TS scheme to continue to recover almost all of a 1.53-dB shaping gain when a turbo code is used in the underlying trellis-code structure.

The coding strategies developed in this paper may find applications in many multiuser system designs where interference presubtraction is a key strategy in achieving the capacity. These applications are expected to include not only the multiuser cancellation of crosstalk in DSL applications mentioned earlier, but also multiuser multiantenna wireless communication systems and high-speed serial-link designs.

## REFERENCES

- [1] M. Costa, "Writing on dirty paper," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 3, pp. 439–441, May 1983.
- [2] G. Ginis and J. M. Cioffi, "Vectored transmission for digital subscriber line systems," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 6, pp. 1085–1104, Jun. 2002.
- [3] G. Caire and S. Shamai, "On the achievable throughput of a multi-antenna Gaussian broadcast channel," *IEEE Trans. Inf. Theory*, vol. 49, no. 7, pp. 1691–1706, Jul. 2003.
- [4] W. Yu and J. M. Cioffi, "Sum capacity of Gaussian vector broadcast channels," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 1875–1892, Sep. 2004.
- [5] S. Vishwanath, N. Jindal, and A. Goldsmith, "Duality, achievable rates, and sum-rate capacity of Gaussian MIMO broadcast channels," *IEEE Trans. Inf. Theory*, vol. 49, no. 10, pp. 2658–2668, Oct. 2003.
- [6] P. Viswanath and D. Tse, "Sum capacity of the multiple-antenna Gaussian broadcast channel and uplink-downlink duality," *IEEE Trans. Inf. Theory*, vol. 49, no. 7, pp. 1912–1921, Jul. 2003.
- [7] H. Weingarten, Y. Steinberg, and S. Shamai, "The capacity region of the Gaussian MIMO broadcast channels," *IEEE Trans. Inf. Theory*, submitted for publication.
- [8] M. Tomlinson, "New automatic equalizer employing modulo arithmetic," *Electron. Lett.*, vol. 7, pp. 138–139, Mar. 1971.
- [9] M. 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, Jun. 1969.
- [10] S. Shamai and R. Laroia, "The intersymbol interference channel: Lower bounds on capacity and channel precoding loss," *IEEE Trans. Inf. Theory*, vol. 42, no. 9, pp. 1388–1404, Sep. 1996.
- [11] U. Erez, S. Shamai, and R. Zamir, "Capacity and lattice strategies for cancelling known interference," in *Proc. Int. Symp. Inf. Theory Applicat.*, Nov. 2000, pp. 681–684.
- [12] G. D. Forney, Jr., "Trellis shaping," *IEEE Trans. Inf. Theory*, vol. 38, no. 3, pp. 281–300, Mar. 1992.
- [13] M. V. Eyuboglu and G. D. Forney, Jr., "Trellis precoding: Combined coding, precoding and shaping for intersymbol interference channels," *IEEE Trans. Inf. Theory*, vol. 38, no. 3, pp. 301–314, Mar. 1992.
- [14] G. Caire and S. Shamai, "Writing on dirty tape with LDPC codes," in *Multiantenna Channels: Capacity, Coding and Signal Processing*, ser. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. Providence, RI: American Mathematics Society, 2003, vol. 62, pp. 123–140.
- [15] R. Zamir, S. Shamai, and U. Erez, "Nested linear/lattice codes for structured multiterminal binning," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1250–1276, Jun. 2002.
- [16] U. Erez and S. ten Brink, "Approaching the dirty-paper limit for cancelling known interference," in *Proc. Allerton Conf. Commun., Control, Comput.*, Oct. 2003, pp. 799–808.
- [17] R. D. Wesel and J. M. Cioffi, "Achievable rates for Tomlinson–Harashima precoding," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, pp. 824–831, Mar. 1998.
- [18] J. H. Conway and N. J. A. Sloane, "A fast encoding method for lattice codes and quantizers," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 11, pp. 820–824, Nov. 1983.
- [19] D. D. Forney, Jr., "Multidimensional constellations—Part II: Voronoi constellations," *IEEE J. Sel. Areas Commun.*, vol. 7, no. 6, pp. 941–958, Aug. 1989.
- [20] T. R. Fischer, M. W. Marcellin, and M. Wang, "Trellis-coded vector quantization," *IEEE Trans. Inf. Theory*, no. 11, pp. 1551–1566, Nov. 1991.
- [21] R. F. H. Fischer, W. H. Gerstacker, and J. B. Huber, "Dynamics limited precoding, shaping, and blind equalization for fast digital transmission over twisted pair lines," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 12, pp. 1622–1633, Dec. 1995.
- [22] R. F. H. Fischer, C. Stierstorfer, and C. Windpassinger, "Precoding and signal shaping in space–time transmission," in *Proc. 8th Canadian Workshop Inf. Theory*, 2003, pp. 83–87.
- [23] —, "Modulo-lattice reduction in precoding schemes," in *Proc. IEEE Int. Symp. Inf. Theory*, 2003, p. 163.
- [24] F. R. Kschischang, H.-A. Loeliger, and B. J. Frey, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.



**Wei Yu** (S'97-M'02) received the B.A.Sc. degree in computer engineering and mathematics from the University of Waterloo, Waterloo, ON, Canada, in 1997, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1998 and 2002, respectively.

Since 2002, he has been an Assistant Professor with the Electrical and Computer Engineering Department at the University of Toronto, Toronto, ON, Canada, where he also holds a Canada Research Chair. His main research interests include multiuser

information theory, optimization, wireless communications, and broadband access networks.



**David P. Varodayan** received the B.A.Sc. degree in engineering science from the University of Toronto, Toronto, ON, Canada, in 2003. He is currently working toward the M.S. and Ph.D. degrees in electrical engineering at Stanford University, Stanford, CA.

His research interests include dirty-paper coding and distributed source coding.



**John M. Cioffi** (S'77–M'78–SM'90–F'96) received the B.S. degree from the University of Illinois, Urbana-Champaign, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984, both in electrical engineering.

He was with Bell Laboratories from 1978 to 1984, and IBM Research from 1984 to 1986. He has been a Professor of Electrical Engineering with Stanford University since 1986. He founded Amati Communications Corporation in 1991 (purchased by Texas Instruments in 1997), and was Officer/Director from 1991 to 1997. He currently is on the Board of Directors of Marvell, Teknovus, ASSIA, Teranetics, and ClariPhy. He is on the advisory boards of Halisos Networks, Ikanos, and Portview Ventures. His specific interests are in the area of high-performance digital transmission. He has published over 250 papers and holds over 40 patents.

Dr. Cioffi is a member of the National Academy of Engineering. He was the recipient of the Hitachi America Professorship in Electrical Engineering at Stanford (2002), the IEEE Kobayashi Medal (2001), the IEEE Millennium Medal (2000), the IEE J. J. Tomson Medal (2000), the 1999 University of Illinois Outstanding Alumnus Award, the 1991 IEEE Communications Magazine Best Paper Award, the 1995 ANSI T1 Outstanding Achievement Award, and the National Science Foundation Presidential Investigator Award (1987–1992).