Published in IET Communications
Received on 30th July 2008
Revised on 13th February 2009
doi: 10.1049/iet-com.2008.0453

ISSN 1751-8628

Compound precoding: a pre-equalisation technique for the bandlimited Gaussian channel

M.F. Flanagan[1] *M. McLaughlin*[2] *A.D. Fagan*[1]
[1]*Department of Electronic and Electrical Engineering, University College Dublin, Belfield, Dublin 4, Ireland*
[2]*DecaWave Ltd., Digital Depot, Thomas St., Dublin 8, Ireland*
E-mail: mark.flanagan@ieee.org

Abstract: Compound precoding is a state-of-the-art technique for combining trellis coding and decision-feedback equalisation (DFE) prior to upstream transmission on the telephone-line channel, and is an option in the International Telecommunications Union (ITU)-T V.92 standard. In this paper, we provide a detailed overview of this technique. We demonstrate that compound precoding combines in a straightforward manner with most practical trellis codes. We show that compound precoding exhibits a signal-to-noise ratio (SNR) gain with respect to other schemes which combine with trellis coding, and quantify this gain. We also show that, for stability of the compound precoder, the precoder feedforward filter must be decomposed into its constituent minimum phase (MP) and all pass (AP) components. Finally, a simulation study of V.92 upstream transmission demonstrates the shaping advantage of compound precoding over competitor techniques for pre-equalisation in this setting.

1 Introduction

In recent years, high-speed data transmission on a bandlimited Gaussian channel (such as may model the telephone channel) has approached closely the theoretical limits predicted by Shannon [1]. This is due mainly to two milestones in the literature. Firstly, the introduction of trellis-coded modulation [2] by Ungerboeck meant that channel coding gains of 3–6 dB could be achieved without bandwidth expansion, and secondly the development of equalisation techniques such as decision-feedback equalisation (DFE) [3] allowed very high signalling rates to be achieved without ISI signal degradation. The combination of DFE with trellis coding presents a dilemma however; when making a symbol decision the DFE requires that all previous symbol decisions are available for feedback, but decoding of the trellis code necessarily implies that these decisions are only available after some delay. Tomlinson-Harashima (TH) precoding [4, 5], on the other hand, has been shown to provide ISI cancellation in a manner transparent to most practically useful trellis codes, and is capable of realising the performance of the ideal zero-forcing DFE (ZF-DFE) by implementing the feedback filter of the equaliser in the transmitter.

The V.34 modem standard [6] incorporated high-rate signalling so as to utilise all of the available bandwidth. Because of the severe attenuation near the band edges, the minimum mean-squared error DFE (MMSE-DFE) structure was required rather than the ZF-DFE. The dilemma of zero-delay decisions was resolved by relocating the feedback filter of the DFE to the transmitter as in TH precoding. V.34 used a precoding method called ISI coding [7] to achieve a data rate of 28.8 kbps (this was later raised to 33.6 kbps in V.34bis).

The V.90 standard introduced the idea of moving from the classical bandlimited additive white Gaussian noise (AWGN) channel model towards a more realistic model of the telephone-line channel. For upstream (user to internet service provider or ISP) communication it used the same modulation, coding and precoding as V.34; however, in the downstream (ISP to user) direction it achieved a data rate of over 50 kbps by using a subset of the pulse code modulation (PCM) quantisation levels as a pulse amplitude modulation (PAM) constellation. This eliminated the quantisation noise introduced by the downstream digital to analogue converter (DAC), which was the primary channel impairment and was previously (and falsely) treated as

AWGN. The symbol rate used was 8000 symbols per second. The V.92 standard [8] extended this scheme to the more difficult upstream direction. Data rates of over 40 kbps were achieved upstream, again by taking advantage of the PCM connections. Precoding to combine trellis coding and DFE was again an intrinsic feature of this standard; some details are given in [8]. The fundamental philosophy of V.92 is to minimise DAC quantisation noise by (i) one-dimensional (PAM) signalling with a signal constellation equal to a subset of the PCM quantisation levels, and (ii) performing all equalisation prior to transmission. As an alternative to the pre-equalisation structure of V.92, equalisation structures based on banks of MMSE subequalisers have been proposed which can cope with severe channel distortion on the uplink [9, 10].

Compound precoding, the central idea of which was first published in [11], is a generalisation of TH precoding which may realise ideal unbiased MMSE-DFE in a power efficient manner by implementing the equaliser feedforward and feedback filters in the transmitter. A version of compound precoding is an option in the V.92 standard [12–14]. This paper presents an overview and analysis of this pre-equalisation technique, which combines trellis coding and DFE for upstream transmission on the telephone-line channel, demonstrating its compatibility with trellis coding. It is shown that compound precoding exhibits a signal-to-noise ratio (SNR) gain with respect to other schemes which combine DFE with trellis coding, and this gain is quantified. It is also shown that, for stability of the compound precoder, the precoder feedforward filter must be decomposed into its constituent minimum phase (MP) and all pass (AP) components. Also, simulations for V.92 transmissions demonstrate the efficacy of this pre-equalisation scheme.

The paper is organised as follows. Section 2 briefly reviews trellis coding and DFE. Section 3 introduces and analyses the compound precoding scheme. Section 4 presents simulation results for V.92 with compound precoding and Section 5 concludes this work.

2 Trellis coding and DFE

In this section we provide a brief background on Ungerboeck trellis codes from the coset code point of view [15]. We then review the result of [16, 17] which identifies the unbiased MMSE-DFE as a canonical structure for equalisation at any SNR when using trellis coding.

2.1 Preliminaries: lattices and cosets

An N-dimensional lattice Λ is simply a subgroup of $\mathbb{R}^N$ (with regard to ordinary vector addition) containing a countable number of elements. Such a subgroup may be imagined as a regular array of points in N-dimensional space. Two points $x, y \in \mathbb{R}^N$ are said to be congruent modulo Λ if and only if $x - y \in \Lambda$; this relation may be written $x \equiv y \pmod{\Lambda}$. If $\Lambda = M\mathbb{Z}$ for some $M \in \mathbb{Z}$, we may use the more familiar notation $x \equiv y \pmod{M}$. A sublattice Λ' of a lattice Λ is a subgroup of Λ. A coset of Λ' in Λ is a set

$$\Lambda' + c = \{\lambda + c | \lambda \in \Lambda'\}$$

where $c \in \Lambda$ is a fixed point which specifies the coset. From elementary group theory, the sublattice Λ' induces a partition of Λ, denoted Λ/Λ', into a disjoint union of equivalence classes such that

- Any two points from the same class are congruent modulo Λ'.
- Any two points from different classes are not congruent modulo Λ'.

These equivalence classes are the distinct cosets of Λ' in Λ; the number of these cosets is denoted Λ/Λ' and is called the *order* of the partition Λ/Λ'.

2.2 Coset codes and trellis codes

An encoder for a coset code is shown in Fig. 1. This encoder is based on a lattice/sublattice partition Λ/Λ' of order 2^{n_c}. Here m information bits are to be encoded per N-dimensional symbol. The k_c-bit block $\boldsymbol{v}_k$ forms the input to a rate k_c/n_c binary encoder (for the block or convolutional code C) to produce the n_c-bit block $\boldsymbol{f}_k$, where $n_c = k_c + r_c$. These n_c bits then select a coset from the partition Λ/Λ'. The remaining $m - k_c$ information bits $\boldsymbol{p}_k$ choose a signal point c_k, lying in this coset, from some finite subset A of Λ called the signal constellation. [This is not entirely true; due to power efficiency considerations, the actual signal constellation is a translate of this subset. However, this fact does not impinge on the present discussion.] This constellation consists of all points in Λ that lie inside some bounded region R. Of course the actual transmit sequence is obtained by demultiplexing the N-dimensional output symbol stream $\{c_k\}$ into constituent PAM or QAM

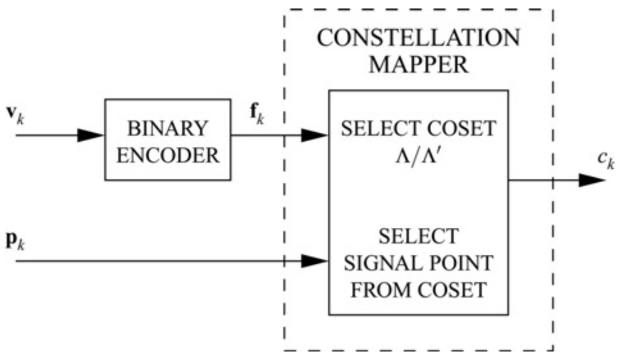

Figure 1 *Encoder for a coset code. A special case of this encoder is an Ungerboeck trellis encoder where the binary encoder is a convolutional encoder, the lattice Λ is $\mathbb{Z}^N$, and $r_c = 1$*

symbols. The coset code $\bar{C}$ is defined to be the set of all signal point sequences $\{c_k\}$ which lie in a sequence of cosets which could be specified by a sequence of coded bits from C. The main point is that only certain sequences of cosets are allowed. If the binary code is a block code, the coset code is called a lattice code; if the binary code is convolutional, the coset code is called a trellis code. We shall consider only the Ungerboeck trellis codes for which $\Lambda = \mathbb{Z}^N$ and $r_c = 1$. [There are also Ungerboeck codes designed for PSK modulation; these are coset codes based on groups which are not lattices.]

The maximum-likelihood (ML) path through the trellis of a coset code may be calculated using the Viterbi Algorithm (VA). It is easy to see that there are 2^{m-k_c} parallel branches in the trellis for each state transition; therefore, the first stage of the VA consists of making a decision on the most likely signal point in the constellation within each coset, and storing its metric ('subset decoding'). The coding gain of a trellis code $\bar{C}$ is given by

$$\gamma_c(\bar{C}) = d^2_{\text{free}}(\bar{C})2^{-2r_c/N} \quad (1)$$

where $d_{\text{free}}(\bar{C})$ denotes the 'free distance' of the trellis code, i.e. the minimum Euclidean distance between signal sequences in $\bar{C}$. This expression for coding gain assumes that the constellation bounding region R is an N-dimensional cube. Of course the only restrictions on the region R are (i) that it must contain 2^{m+r_c} signal points, and (ii) that it must contain an equal number of signal points from each coset. Therefore R need not be a cubical region, and in fact a different choice of constellation boundary may give a further gain – this is called 'constellation shaping' [18, 19]. The code and constellation mapping are designed jointly in order to maximise $d^2_{\text{free}}(\bar{C})$ [20].

2.3 DFE

The structure of the decision-feedback equaliser (DFE) is shown in Fig. 2, where we assume the minimum mean-square error solution for the filter coefficients (MMSE-DFE). In an important two-part paper Cioffi *et al.* [16, 17] generalised Price's result on the ZF-DFE at high SNR [21] to the unbiased MMSE-DFE at arbitrary SNR, thus identifying the unbiased MMSE-DFE as a canonical structure for use with coded modulation.

Assuming this DFE is ideal ($\hat{c}_k = c_k \forall k$), and also that the filters have infinite length, Cioffi *et al.* showed that this

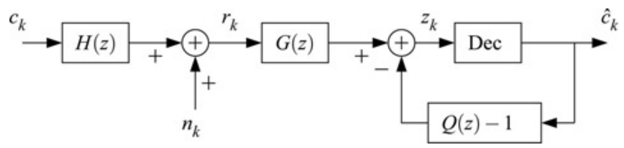

Figure 2 *Structure of the ideal MMSE-DFE. The ideal assumption ensures* $\hat{c}_k = c_k \ \forall k$

receiver is *biased* [16]. This bias may be removed by inserting a scaling factor β immediately prior to the decision device. (At high SNR this scaling factor is approximately unity and may be ignored.) The resulting receiver, called the unbiased MMSE-DFE (MMSE-DFE,U), has higher MSE but better error performance. It was also shown in [16] that this receiver is in fact the MMSE-optimum receiver over all unbiased DFE structures.

The SNR of the unbiased MMSE-DFE receiver is defined as [16]

$$\text{SNR}_{\text{MMSE-DFE,U}} \triangleq \frac{E\{c_k^2\}}{E\{(\beta z_k - c_k)^2\}}$$

It was shown in [17] that the channel capacity (in bits per symbol) is related to the SNR of the ideal unbiased MMSE-DFE by

$$C(T) = \frac{1}{2}\log_2(1 + \text{SNR}_{\text{MMSE-DFE,U}})$$

This formula generalises Price's result [21] on an approximate relation between channel capacity and SNR of the ZF-DFE at high SNR to an *exact* relation between channel capacity and SNR of the unbiased MMSE-DFE at *any* SNR. Price's result may be taken as a high-SNR special case of this result, as at high SNR the MMSE-DFE tends to the ZF-DFE.

It is also shown in [17] that, when uncoded PAM is used with an unbiased MMSE-DFE, the gap to capacity for an arbitrary bandlimited AWGN channel is approximately independent of the channel characteristics for all values of $\text{SNR}_{\text{MMSE-DFE,U}}$, provided the assumption of ideal DFE holds true. Of course relocating the feedback filter to the transmitter as in TH precoding will maintain this condition. Finally, Cioffi *et al.* [17] argue that, for a particular target symbol error probability, trellis coding and shaping techniques yield approximately the same gain on any arbitrary bandlimited AWGN channel with MMSE-DFE,U as on the ideal channel, and thus trellis coding, shaping and ideal MMSE-DFE,U (through precoding) may close the gap to capacity to the same extent on any strictly bandlimited AWGN channel.

3 Compound precoding

In this section, we introduce the compound precoding scheme as a power efficient means of implementing the ideal unbiased MMSE-DFE in the transmitter and explain why this is an effective method to combat quantisation noise in upstream V.92 transmission. We show that in order to ensure the stability of the compound precoder, it is necessary to decompose the precoder feedforward filter $G(z)$ into its constituent MP and AP components. We demonstrate that

IET Commun., 2009, Vol. 3, Iss. 8, pp. 1432–1441
doi: 10.1049/iet-com.2008.0453

compound precoding combines straightforwardly with most Ungerboeck trellis codes. We then describe two other methods of achieving ideal unbiased MMSE-DFE through precoding, and show that, under quite general conditions, the compound precoding scheme yields an SNR gain over these methods, which we quantify. We also mention another advantage of compound precoding arising from constellation shaping considerations.

3.1 TH precoding

The TH precoding scheme [4, 5] is shown in Fig. 3. The input symbols c_k are assumed to be uncoded M-ary PAM symbols, i.e. $c_k \in \{\pm 1/2, \pm 3/2, \pm 5/2, \ldots, \pm (M-1)/2\}$. The modulo devices are understood to perform modulo reduction to the interval $(-M/2, M/2]$. Assume that the channel can be modelled as a monic finite impulse response (FIR) filter $H(z)$ with AWGN n_k of variance σ^2. An estimate of the channel coefficients is made at the receiver, and this estimate is sent back to the transmitter. This gives the filter $Q(z) - 1$ shown in the figure. Analysis of this system is straightforward; if we denote the Z-domain transmit signal as $T(z)$, then

$$T(z) = C(z) - MS(z) - [Q(z) - 1]T(z)$$

the subtraction of $MS(z)$ being due to the modulo operation ($s_k \in \mathbb{Z}$ for each k). The received sequence has Z-transform

$$R(z) = H(z)T(z) + N(z)$$

Assuming perfect channel estimation, $Q(z) = H(z)$ and so

$$R(z) = C(z) - MS(z) + N(z) \qquad (2)$$

so upon modulo-M reduction, the input to the decision device is

$$z_k \equiv c_k + n_k \qquad (\text{mod } M)$$

Hence, the decision device operates by making a nearest-neighbour decision modulo M. The linear PAM constellation is effectively replaced by a 'circular' constellation such that every point has two nearest neighbours. This leads to a 'data flipping' problem [this term was coined in [22]); primarily, $c_k = (M-1)/2$ being detected as $c_k = -(M-1)/2$ and vice versa. Hence the probability of symbol error is slightly greater than that for the ideal channel ($H(z) = 1$). This data flipping problem is a direct consequence of the modulo devices. However, if the modulo devices are removed, the precoder will be stable if and only if the channel is MP. As stated in [23], the TH precoding scheme is transparent to many good trellis codes if M is a multiple of 4. In addition to this benefit, TH precoding also eliminates the problem of error propagation which would be a possibility if the recursive filter $Q(z) - 1$ were implemented in the receiver (ZF-DFE). The cost of TH precoding is the need to send the channel estimates back to the transmitter, a loss in the tracking capability of the channel estimation algorithm, and a power increase of approximately $10\log_{10}(M^2/(M^2-1))$ dB. It is found in practice that for general channel filters $H(z)$, the elements of the sequence $\{t_k\}$ are approximately independent and uniformly distributed on $(-M/2, M/2]$ [23]. The TH precoder may be trained at the receiver as would a ZF-DFE.

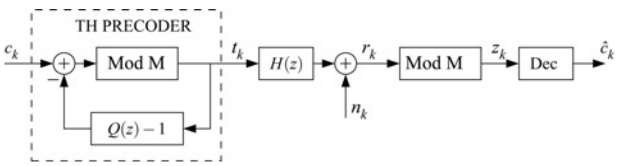

Figure 3 *TH precoding*

3.2 Compound precoding scheme

The discrete-time model for a system employing compound precoding is shown in Fig. 4. The compound precoder is a precoder which incorporates both feedforward and feedback equaliser filters, and thus may accommodate ideal unbiased MMSE-DFE. Recall that a primary impediment to reliable communication over this channel is the quantisation noise introduced by the analogue to digital converter (PCM quantiser) at the telephone exchange. The additive noise n_k in Fig. 4 models the sum of this quantisation noise and channel output AWGN. The fundamental philosophy of V.92 is to minimize this quantisation noise by (i) one-dimensional (PAM) signalling with a signal constellation equal to a subset of the PCM quantisation levels, and (ii) performing all equalisation prior to transmission. For adequate equalisation at the transmitter, both feedforward and feedback filters are required, in general – of course the unbiased MMSE-DFE is a canonical choice in this regard. The compound precoder may implement the ideal unbiased MMSE-DFE structure through precoding in a power efficient manner. In order to simplify the analysis which follows, we will assume that the PCM quantisation levels are uniformly spaced, and thus we may assume that the input symbol c_k lies in a uniform M-ary PAM constellation, i.e. $c_k \in \{\pm 1/2, \pm 3/2, \pm 5/2, \ldots, \pm (M-1)/2\}$. This assumption is also made in [11–13]. Note however that in a practical implementation the PCM levels lie on an A-law (Europe) or μ-law (USA) characteristic.

The structure of the compound precoder is shown in Fig. 5. The filters $G(z)$ and $Q(z) - 1$ are estimates of the feedforward and feedback filters, respectively, of an unbiased MMSE-DFE; these are trained at the receiver and then sent back to the transmitter (cf. Fig. 2). Assume for the moment that the input sequence $\{c_k\}$ is uncoded. From the input symbol c_k is subtracted Ms_k where $s_k \in \mathbb{Z}$

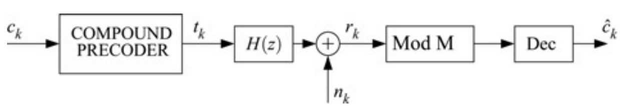

Figure 4 *Discrete-time model for a system using compound precoding*

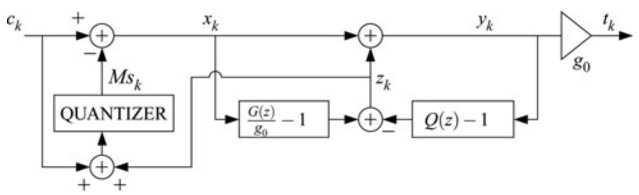

Figure 5 *Full structure of the compound precoder*

for each k. The values of the $\{s_k\}$ are chosen so as to constrain the signal y_k to the interval $(-M/2, M/2]$, thus preventing instability of the precoder; this is possible since $G(z)/g_0$ is monic. The quantiser of Fig. 5 quantises its input, $c_k + z_k$, to the nearest multiple of M (denoted Ms_k) such that the quantisation error $m_k = Ms_k - (c_k + z_k)$ lies in $[-M/2, M/2)$. To show that the correct integer s_k is computed, observe that

$$y_k = c_k - Ms_k + z_k$$

Thus we have $y_k = -m_k$ and so $y_k \in (-M/2, M/2]$. Thus the stability of the precoder output sequence is assured. Similar to the case of TH precoding, in practice the elements of the sequence $\{y_k\}$ are found to be approximately independent and uniformly distributed on $(-M/2, M/2]$.

Similarly to the case of uncoded TH precoded PAM, if we assume that the equalisation filters approximately cancel the channel filter $(Q(z)/G(z) \simeq H(z))$, then (from Fig. 5) the received symbol is simply $r_k = c_k - Ms_k + n_k$, where n_k is AWGN. Since $c_k \in \{\pm 1/2, \pm 3/2, \pm 5/2, \ldots, \pm(M-1)/2\}$, the term Ms_k may be removed by a modulo device at the receiver which reduces its input to the interval $(-M/2, M/2]$. This system provides almost the same performance as ideal unbiased MMSE-DFE, except that it exhibits the same 'data flipping' problem as the TH precoder.

3.3 Stability of the compound precoder

The compound precoder of Fig. 5 is not always stable. What is meant by this statement is that even though the sequences $\{c_k\}$, $\{y_k\}$ and $\{m_k\}$ are bounded, the sequences $\{z_k\}$, $\{s_k\}$ and $\{x_k\}$ need not be because of the feedback loop. Observe from Fig. 5 that

$$Z(z) = \left[\frac{G(z)}{g_0} - 1\right]X(z) - [Q(z) - 1][X(z) + Z(z)] \quad (3)$$

and

$$X(z) = -[M(z) + Z(z)] \quad (4)$$

Solving (3) and (4) simultaneously yields

$$X(z) = -\frac{g_0 Q(z) M(z)}{G(z)}$$

Since $\{m_k\}$ is bounded, the compound precoder circuit is stable if and only if the feedforward filter $G(z)$ is MP (note that the boundedness of $\{x_k\}$ implies that of $\{z_k\}$ and $\{s_k\}$). If $G(z)$ is not MP, it may be factored [24] into a MP component $W(z)$ and an AP component $C(z)/C^B(z)$ (Here $C^B(z)$ denotes the reciprocal filter of $C(z)$, which is obtained by reversing the coefficients of $C(z)$, i.e. $C^B(z) = c_p + c_{p-1}z^{-1} + \cdots + c_0 z^{-p}$ if $C(z) = c_0 + c_1 z^{-1} + \cdots + c_p z^{-p}$.), i.e.

$$G(z) = W(z)\frac{C(z)}{C^B(z)} \quad (5)$$

For implementation purposes, we generally take $C^B(z)$ to be monic. The compound precoder is then constructed as illustrated in Fig. 6. Again, the elements of the sequence $\{y_k\}$ are found to be approximately independent and uniformly distributed on $(-M/2, M/2]$.

3.4 Combination of compound precoding with trellis coding

The combination of compound precoding with trellis coding is illustrated in Fig. 7. In this combination, there are two main issues to be addressed.

- First, we must ensure that the received sequence is a valid sequence in the original trellis code $\bar{C}$, corrupted by AWGN.
- Second, the use of TH precoding with trellis coding necessitates a slight modification to the trellis decoder (VA) in the receiver.

3.4.1 Coset invariance: If we assume that the equalisation filters approximately cancel the channel filter $(Q(z)/G(z) \simeq H(z))$, then it may be easily seen from the compound precoder model of Fig. 5 that the net effect of the (compound precoder, channel filter) combination is that of replacing each N-dimensional signal point c_k with the point $a_k = c_k - x_k$, where $x_k \in M\mathbb{Z}^N$. The issue for trellis decoding here is whether the signal point sequence $\{a_k\}$ is in $\bar{C}$. We shall now prove that for all of the lattice-type coset codes listed in [25], this is true provided M is a multiple of 4,

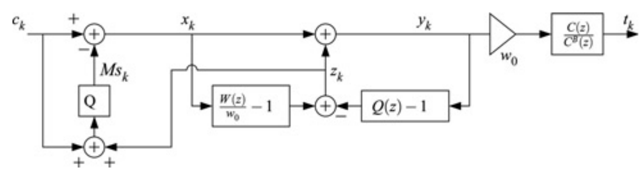

Figure 6 *Stable compound precoder assuming non-minimum phase feedforward filter G(z). This circuit uses the MP–AP decomposition of G(z)*

Figure 7 *System employing compound precoding in conjunction with trellis coding*

doi: 10.1049/iet-com.2008.0453

and furthermore, that the signal point a_k lies in the same coset $\Lambda' + c$ as the original point c_k. The proofs follow for each dimensionality of trellis code.

One-dimensional code ($N = 1$): Here the lattice $\Lambda = \mathbb{Z}$, and $k_c = r_c = 1$. The sublattice used is $\Lambda' = 4\mathbb{Z}$, and thus $x_k \in \Lambda'$, since M is a multiple of 4. It follows that $a_k = c_k - x_k \in \Lambda' + c$.

Two-dimensional code ($N = 2$): Here the lattice $\Lambda = \mathbb{Z}^2$. The maximum k_c used is 2 and $r_c = 1$. The sublattice used in this case is $\Lambda' = 2\boldsymbol{R}\mathbb{Z}^2$, where $\boldsymbol{R} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. This sublattice may be written $2\boldsymbol{R}\mathbb{Z}^2 = \left\{ 2\begin{pmatrix} a+b \\ a-b \end{pmatrix} | a, b \in \mathbb{Z} \right\}$. Now, since M is a multiple of 4, we may write $x_k = 4\begin{pmatrix} m \\ n \end{pmatrix}$, for some integers m and n. Therefore

$$x_k = 2\begin{pmatrix} 2m \\ 2n \end{pmatrix} = 2\begin{pmatrix} (m+n)+(m-n) \\ (m+n)-(m-n) \end{pmatrix}$$

and thus $x_k \in \Lambda'$. It follows that $a_k = c_k - x_k \in \Lambda' + c$.

Four-dimensional code ($N = 4$): Here the lattice $\Lambda = \mathbb{Z}^4$. The maximum k_c used is 4 and $r_c = 1$. The sublattice used in this case is $\Lambda' = 2D_4$, where

$$D_4 = \left\{ (a \quad b \quad c \quad d)^{\mathrm{T}} | a + b + c + d \equiv 0 \pmod{2} \right\}$$

Since M is a multiple of 4, we may write $x_k = 4\,(x \;\; y \;\; z \;\; w)^{\mathrm{T}}$, for some integers x, y, z, w. Therefore

$$x_k = 2(2x \quad 2y \quad 2z \quad 2w)^{\mathrm{T}} \in 2D_4$$

Thus $x_k \in \Lambda'$, and it follows that $a_k = c_k - x_k \in \Lambda' + c$.

Eight-dimensional code ($N = 8$): Here the lattice $\Lambda = \mathbb{Z}^8$. The maximum k_c used is again 4 and $r_c = 1$. The sublattice used in this case is $\Lambda' = \boldsymbol{R}'D_8$, where

$$\boldsymbol{R}' = \begin{pmatrix} \boldsymbol{R} & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{R} & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{R} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{R} \end{pmatrix}$$

and $D_8 = \{(a_1 \quad a_2 \quad \cdots \quad a_8)^{\mathrm{T}} | \sum_i a_i \equiv 0 \pmod{2}\}$. If, for each $i \in \{1, 2, 3, 4\}$, we define $\boldsymbol{a}_i = \begin{pmatrix} a_{2i-1} \\ a_{2i} \end{pmatrix}$, the sublattice may be written

$$\boldsymbol{R}'D_8 = \left\{ \begin{pmatrix} \boldsymbol{R}\boldsymbol{a}_1 \\ \boldsymbol{R}\boldsymbol{a}_2 \\ \boldsymbol{R}\boldsymbol{a}_3 \\ \boldsymbol{R}\boldsymbol{a}_4 \end{pmatrix} | \sum_i a_i \equiv 0 \pmod{2} \right\} \qquad (6)$$

Now, since M is a multiple of 4, we may write $x_k = 4\left(\boldsymbol{n}_1^{\mathrm{T}} \quad \boldsymbol{n}_2^{\mathrm{T}} \quad \boldsymbol{n}_3^{\mathrm{T}} \quad \boldsymbol{n}_4^{\mathrm{T}}\right)^{\mathrm{T}}$, where $\boldsymbol{n}_i = \begin{pmatrix} n_{2i-1} \\ n_{2i} \end{pmatrix} \in \mathbb{Z}^2$ for each $i = 1, 2, 3, 4$. Now

$$4\boldsymbol{n}_i = \begin{pmatrix} 4n_{2i-1} \\ 4n_{2i} \end{pmatrix} = \boldsymbol{R}\begin{pmatrix} 2(n_{2i-1} + n_{2i}) \\ 2(n_{2i-1} - n_{2i}) \end{pmatrix}$$

and thus, by (6), x_k is an element of $\boldsymbol{R}'D_8$, where we let $\boldsymbol{a}_i = \begin{pmatrix} 2(n_{2i-1} + n_{2i}) \\ 2(n_{2i-1} - n_{2i}) \end{pmatrix}$ for each $i = 1, 2, 3, 4$. Thus $x_k \in \Lambda'$, and it follows that $a_k = c_k - x_k \in \Lambda' + c$.

Note: In each case, we have considered only the code with maximum k_c, for which the partition is $\mathbb{Z}^N/\Lambda'$. Any trellis code with a lower value of k_c uses a sublattice Λ'' which contains Λ' as a sublattice. Each proof here establishes that $x_k \in \Lambda'$, and thus $x_k \in \Lambda''$ also. Thus the proofs for all the codes of [25] with lower values of k_c follow trivially.

3.4.2 VA modification: Since we have established that the received sequence is indeed a valid sequence in the trellis code $\bar{C}$ perturbed by AWGN, the VA-based trellis decoder may be used for the detection of the signal point sequence $\{a_k\}$ in the same way that it is normally used for detection of the signal point sequence $\{c_k\}$. The only difference is that the modified signal point $a_k = c_k - x_k$ is no longer restricted to the finite signal constellation within the bounded region R, but may be any sequence in $\bar{C}$. Therefore the VA needs to be modified slightly. The first part of Viterbi trellis decoding, called subset decoding, usually involves choosing the nearest point *within the signal constellation*, in a particular coset, to the received point r_k, and storing its metric; it now involves choosing the nearest point *in the entire coset* to the received point r_k, and storing its metric [22]. The subsequent processing of the trellis decoder is as before. In this case the VA no longer performs ML sequence detection, but is a good approximation to a ML sequence detector [7]. When the sequence $\{\hat{a}_k\}$ has been detected, a modulo device reduces this sequence to the corresponding sequence of constellation symbols $\{\hat{c}_k\}$. Finally, the constellation symbol estimates are demapped to produce information bit estimates $\{\hat{u}_k\}$. Since the coding gain of the trellis code is independent of the bounding region R [by (1)], this scheme allows the full equalisation gain of the compound precoder and the full coding gain of the trellis code to be realised jointly.

3.5 SNR gain of compound precoding

In this section we show that, with or without trellis coding, the compound precoding scheme exhibits an SNR gain with respect to other TH-based precoding schemes which incorporate MMSE-DFE, and we give a concise expression for this SNR gain. The two other schemes in question are: TH precoding followed by linear post-equalisation and TH precoding with linear pre-equalisation (LPE).

3.5.1 TH precoding followed by linear post-equalisation: In this scheme, the feedforward and feedback filters of an MMSE-DFE are trained at the receiver, then the feedback filter $Q(z) - 1$ is sent back to the transmitter. This filter is incorporated into the TH precoder in the usual way. The required feedforward equalisation is performed at the receiver by filtering the received sequence by $G(z)$. Note that while we include this scheme in our analysis, it is not a suitable scheme for V.92 as all equalisation is not performed prior to transmission.

3.5.2 TH precoding with LPE: In this scheme, the feedforward and feedback filters of an MMSE-DFE, again trained at the receiver, are both sent back to the transmitter. This time, after standard TH precoding [which again incorporates the filter $Q(z) - 1$], the signal is filtered by the feedforward filter $G(z)$ prior to transmission. Thus all equalisation is performed prior to transmission.

3.5.3 Compound precoding: In this scheme, both the feedforward filter and the feedback filter of the MMSE-DFE are incorporated into the precoder, so there is no need for pre-filtering or post-filtering.

3.5.4 Comparison of methods: In order to make a fair comparison between the three schemes, we make a few initial assumptions. We assume that all three schemes transmit PAM symbols at the same rate $1/T$ over the same continuous-time channel; hence the discrete-time channel model is the same in all three cases. We assume without loss of generality that all three schemes use trellis coding; the case of uncoded M-ary PAM may be catered for within this framework by taking $\bar{C}$ to be the universe code consisting of all M-ary PAM sequences. Moreover, we assume that all three schemes use the same N-dimensional trellis code $\bar{C}$. Therefore the three schemes all transmit information at the same rate. Finally, we assume that the equaliser filters approximately cancel the channel filter, i.e. $Q(z)/G(z) \simeq H(z)$.

In the two schemes using TH precoding, the output of the TH precoder, denoted y_k in both cases, is approximately uniformly distributed on the interval $(-M/2, M/2]$; we denote its average energy by $P\,(= M^2/12)$. In the case of TH precoding with LPE, the received sequence $\{r_k\}$ is approximately a sequence in $\bar{C}$ corrupted by AWGN of variance σ^2 per dimension; however, the transmit symbol here, t_k, has an average energy approximately equal to $P\sum_{i=0}^{p} g_i^2$, since the sequence $\{y_k\}$ is approximately uncorrelated. In the case of TH precoding followed by linear post-equalisation, the received sequence after post-equalisation, $\{\rho_k\}$, is approximately a sequence in $\bar{C}$ corrupted by AWGN. The AWGN has variance $\sigma^2\sum_{i=0}^{p} g_i^2$ however, since it passes through the filter $G(z)$. Thus the overall performance of pre-equalised and post-equalised TH precoding schemes are approximately equal. Compound precoding involves no post-filtering of the noise, and so the received sequence r_k in this case is approximately a sequence in $\bar{C}$ corrupted by AWGN of variance σ^2 per dimension. Also, the signal y_k in Fig. 6 is approximately uniformly distributed on $(-M/2, M/2]$; thus the average transmit symbol energy of the system is approximately Pw_0^2. Therefore the system using compound precoding exhibits an SNR gain in dB of

$$\gamma_{CP} = 10\log_{10}\left(\frac{w_0^2}{\sum_{i=0}^{p} g_i^2}\right) = 10\log_{10}\left(\frac{w_0^2}{\sum_{i=0}^{p} w_i^2}\right)$$

with respect to the other two schemes.

We conclude that, under quite general conditions, compound precoding yields an SNR gain which is a function of the MP component of the precoder feedforward filter. However, it is necessary to decompose the precoder feedforward filter into its MP and AP components in order to realise this SNR gain in a practical system.

3.6 Shaping advantage of compound precoding

In this section we describe another principal advantage of compound precoding arising from shaping gain considerations [11]. Consider the generic circuit described by

$$Y(z) = F(z)(C(z) - MS(z)) \tag{7}$$

This circuit may model either the TH precoder with LPE or the compound precoder. From the constellation symbol sequence $\{c_k\}$ is subtracted a sequence $\{Ms_k\}$ where $s_k \in \mathbb{Z}$ for each k, and the resulting sequence is filtered by the filter $F(z)$ to produce the sequence $\{y_k\}$. The values of s_k are chosen to restrict the values of y_k to the interval $(-M/2, M/2]$. For the TH precoder $F(z) = 1/Q(z)$, and for the compound precoder $F(z) = G(z)/g_0Q(z)$, where $Q(z)/G(z) \simeq H(z)$.

In practical implementations of precoding for V.92, it is desirable to restrict the allowable values of s_k to a small range of values centred on the origin. In this respect the compound precoder outperforms the TH precoder, for the following reason. The telephone-line channel frequency response $H(z)$ is usually approximately flat over most of the band; therefore so is $F(z) = G(z)/g_0Q(z)$ and hence the magnitudes of the required s_k are small. However, the filters $G(z)$ and $Q(z)$ often have zeros in approximately the same position in the z-plane; therefore for the TH precoder $F(z) = 1/Q(z)$ may have high gain at some frequency even though the overall channel response is reasonably flat at this frequency. This can lead to larger values of s_k, which gives a performance loss because the higher levels are farther apart and have a power penalty, the constellation being non-linear.

3.7 Practical issues

One important practical issue is how to effect the decomposition of the precoder feedforward filter into its constituent MP and AP components. In [26] a suite of algorithms is given, each of which performs this

doi: 10.1049/iet-com.2008.0453

decomposition by gradient-based adaptive signal processing techniques. In general, the algorithms of [26] require higher complexity (for a given decomposition accuracy) if the system zeros are positioned close to the unit circle in the complex plane. Thus schemes such as the radial root reduction technique of [27] could be beneficial to compound precoding, and achieve a compromise between equalisation performance and complexity.

Another important issue in practice is that of update of the precoder coefficents at the transmitter. In wired channels such as asymmetric digital subscriber line (ADSL), time variation of the channel is relatively slow. Many pre-equalisation schemes adopt the strategy of keeping track of the measured signal to interference plus noise ratio (SINR) (or packet error rate (PER)) at the receiver, and retrain the precoder once the SINR becomes too low (or the PER becomes too high). In state-of-the-art applications, the processing to perform the required decomposition will be negligible compared to the retrain time of the precoder. In fact, if the adaptive algorithms of [26] are used, then adaptive update of the precoder coefficients would be a possibility assuming the availability of a feedback channel.

4 Simulation results

In this section we provide simulation results, which illustrate the shaping advantage of compound precoding in the context of upstream transmission in V.92. The symbol rate is $R_s = 8$ ksymbols/s in order to match the codec at the central office (CO). The PAM constellation is set to match the set of $N = 256$ μ-law PCM levels, which we denote by $\mathcal{S} = \{a_1, a_2, \ldots, a_N\}$ where the elements are in increasing order. The DFE is comprised of feedforward and feedback filters $G(z) = -1.7542 + 0.7238z^{-1} + 0.4967z^{-2} + 0.3904z^{-3} + 0.3285z^{-4} + 0.0690z^{-5} + 0.0242z^{-6} + -0.4026z^{-7} - 0.0258z^{-8}$ and $Q(z) = 1 - 0.2405z^{-1} - 0.7302z^{-2} - 0.1957z^{-3} - 0.3374z^{-4} + 0.0640z^{-5} + 0.0019z^{-6} + 0.3455z^{-7} + 0.0857z^{-8}$, respectively, these being trained using the least-mean-square (LMS) algorithm. The channel filter is EIA loop 3, $H(z) = -0.5701 - 0.1040z^{-1} + 0.0075z^{-2} - 0.0065z^{-3} + 0.0034z^{-4} - 0.0010z^{-5} + 0.0011z^{-6} + 0.0002z^{-7}$, cascaded with a section with zeros at $z = -0.99$, 1, 1 and poles at $z = -0.8$, 0.9, 0.9, modelling the combination of transformer and codec. In this application, the modulo M operation at the transmitter is performed not on the constellation levels $\{a_i\}$ directly, but on their *indices*. For each $k = 0, 1, \ldots, M-1$, the set $\mathcal{A}_k = \{a_i \in \mathcal{S} : i \equiv k \pmod{M}\} \subset \mathcal{S}$. Specifically, the precoder chooses that point in $\mathcal{A}_k$ which minimises the energy of the transmitted symbol. Trellis coding is not simulated explicitly, but is modelled by (a) a constellation design that chooses the largest constellation that maintains a distance of $D_u = 8$ between its points (after addition of echo), and (b) the assumption that errors in received symbols may be corrected if and only if they are to the nearest-neighbour constellation level.

At the receiver, to the output of the channel filter is added echo and AWGN prior to quantisation by the codec to the nearest quantisation level (element of $\mathcal{S}$). The echo is modelled as an uncorrelated sequence of points uniformly chosen from the downstream constellation, attenuated by transformer hybrid losses which are set at 20 dB. The downstream constellation is designed as the largest constellation $\mathcal{X} \subset \mathcal{S}$ which maintains a distance of $D_d = 16$ between its points. The AWGN noise level is -65.5 dBm. Table 1 illustrates some performance results (operating points) of the system in terms of data rate, transmit power and bit error rate (BER). Fig. 8 shows the histogram of transmit signal amplitudes, both for compound precoding and for TH precoding with LPE, demonstrating the shaping gain of compound precoding. The histogram is illustrated for an information rate equal to 48.3 kbits/s. It may be seen that in the case of compound precoding, the lower-energy constellation levels are chosen with a higher probability.

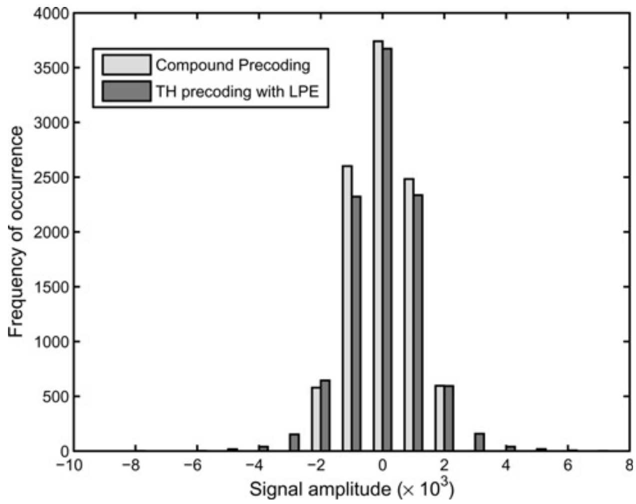

Figure 8 *Transmit signal amplitude distribution for compound precoding, and TH precoding with LPE, illustrating the shaping advantage of compound precoding*

Table 1 Selected performance results of compound precoding for V.92

Transmit power (downstream), dBm	Data rate (downstream), kbps	Transmit power (upstream), dBm	Data rate (upstream), kbps	Bit error rate (upstream), BER
−8.14	54.864	−11.22	45.633	3.75×10^{-4}
−12.19	51.935	−10.99	48.283	5.11×10^{-4}
−16.25	48.000	−10.81	49.934	4.2×10^{-3}

doi: 10.1049/iet-com.2008.0453

5 Conclusion

Compound precoding is a generalization of TH precoding which may realise ideal unbiased MMSE-DFE in a power efficient manner by implementing the equaliser feedforward and feedback filters in the transmitter. A version of compound precoding is an option in the V.92 standard, which requires that all equalisation be performed prior to transmission. We have provided a thorough analysis of this pre-equalisation technique, emphasising its inherent advantages. Also, the necessity of the MP–AP decomposition of the compound precoder feedforward filter to ensure precoder stability was proven, and an expression for the SNR gain of the compound precoder was derived. Simulation results demonstrate the shaping gain with respect to competitive pre-equalisation techniques in the context of upstream transmission in V.92.

6 Acknowledgment

The authors acknowledge the support of Enterprise Ireland and Lake Communications Ltd.

7 References

[1] SHANNON C.E.: 'A mathematical theory of communication', *Bell System Tech. J.*, 1948, **27**, pp. 379–423, 623–656

[2] UNGERBOECK G., CSAJKA I.: 'On improving data-link performance by increasing the channel alphabet and introducing sequence coding'. Int. Symp. Inf. Theory, June 1976

[3] BELFIORE C.A., PARK J.H.: 'Decision-feedback equalization', *Proc. IEEE*, 1979, **67**, pp. 1143–1156

[4] TOMLINSON M.: 'New automatic equalizer employing modulo arithmetic', *Electron. Lett.*, 1971, **7**, pp. 138–139

[5] HARASHIMA H., MIYAKAWA H.: 'Matched-transmission technique for channels with intersymbol interference', *IEEE Trans. Commun.*, 1972, **COM-20**, pp. 774–780

[6] FORNEY G.D., BROWN L., EYUBOGLU M.V., MORAN J.L.: 'The V.34 high-speed modem standard', *IEEE Commun. Mag.*, 1996, **34**, pp. 28–33

[7] LAROIA R.: 'Coding for intersymbol interference channels – combined coding and precoding', *IEEE Trans. Inf. Theory*, 1996, **42**, pp. 1053–1061

[8] KIM D.Y., HUMBLET P.A., EYUBOGLU M.V., BROWN L., FORNEY G.D., MEHRABANZAD S.: 'V.92: the last dial-up modem?', *IEEE Trans. Commun.*, 2004, **52**, pp. 54–61

[9] AYANOGLU E., DAGDEVIREN N.R., GOLDEN G.D., MAZO J.E.: 'An equalizer design technique for the PCM modem: a new modem for the digital public switched network', *IEEE Trans. Commun.*, 1998, **46**, pp. 763–774

[10] LEE J.H., LEE Y.H.: 'Design of multiple MMSE subequalizers for faster-than-nyquist-rate transmission', *IEEE Trans. Commun.*, 2004, **52**, pp. 1257–1264

[11] MCLAUGHLIN M.: 'Compound precoding: pre-equalization for channels with combined feedforward and feedback characteristics'. Irish Signals Syst Conf., 2000, pp. 536–543

[12] Lake Datacomms (McLaughlin M.), 'Adapting the Tomlinson precoding method for channels with combined feedforward and feedback characteristics', Delayed Contribution D385 to ITU-T Study Group 16, February 2000

[13] Motorola Inc. (EYUBOGLU, M. V., FORNEY, G. D., KIM, D., MEHRABANZAD, S.): 'Pre-equalizer structures for recommendation V.92', Delayed Contribution D469 to ITU-T Study Group 16, February 2000

[14] ITU-T Recommendation V.92: 'Enhancements to recommendation V.90', 2000

[15] FORNEY G.D.: 'Coset codes – part I: Introduction and geometrical classification', *IEEE Trans. Inf. Theory*, 1988, **34**, pp. 1123–1151

[16] CIOFFI J.M., DUDEVOIR G.P., EYUBOGLU M.V., FORNEY G.D.: 'MMSE decision-feedback equalizers and coding – part I: Equalization results', *IEEE Trans. Commun.*, 1995, **43**, pp. 2582–2594

[17] CIOFFI J.M., DUDEVOIR G.P., EYUBOGLU M.V., FORNEY G.D.: 'MMSE decision-feedback equalization and coding – part II: Coding results', *IEEE Trans. Commun.*, 1995, **43**, pp. 2595–2604

[18] CALDERBANK A.R., OZAROW L.: 'Nonequiprobable signaling on the Gaussian channel', *IEEE Trans. Inf. Theory*, 1990, **36**, pp. 726–740

[19] FORNEY G.D.: 'Trellis shaping', *IEEE Trans. Inf. Theory*, 1992, **38**, pp. 281–300

[20] UNGERBOECK G.: 'Channel coding with multilevel/phase signals', *IEEE Trans. Inf. Theory*, 1982, **IT-28**, pp. 55–67

[21] PRICE R.: 'Nonlinearly feedback equalized PAM versus capacity for noisy filter channels'. Proc. ICC, June 1972

[22] AMAN A., CUPO R.L., ZERVOS N.A.: 'Combined trellis coding and DFE through Tomlinson precoding', *IEEE J. Sel. Areas Commun.*, 1991, **9**, pp. 876–884

[23] FORNEY G., EYUBOGLU M.V.: 'Combined equalization and coding using precoding', *IEEE Commun. Mag.*, 1991, **29**, (12), pp. 25–34

[24] OPPENHEIM A.V., SCHAFER R.W.: 'Discrete-time signal processing' (Prentice-Hall, Englewood Cliffs, NJ, 1989)

doi: 10.1049/iet-com.2008.0453

[25] UNGERBOECK G.: 'Trellis-coded modulation with redundant signal sets, parts I and II', *IEEE Commun. Mag.*, 1987, **25**, pp. 5–21

[26] FLANAGAN M.F., MCLAUGHLIN M., FAGAN A.D.: 'Gradient-adaptive algorithms for minimum phase – all pass decomposition of a finite impulse response system', *IET Signal Process.*, To be published

[27] SELLARS M.P., GREAVES S.D., WASSELL I.J., PORTER J., FITZGERALD W.J.: 'Stabilized precoder for indoor radio communications', *IEEE Commun. Lett.*, 2000, **4**, pp. 315–317