T1E1.4/96-066

STUDY PROJECT: HIGH SPEED DIGITAL SUBSCRIBER LINES

****

SOURCE:
University of California, Los Angeles
Gregory J. Pottie
Electrical Engineering Department
405 Hilgard Ave.
Los Angeles CA 90095
(310) 825-8150

****

ISSUES ADDRESSED: Preliminary discussion of shaping and precoding for single-pair high speed digital subscriber lines

****

DATE: April 22, 1996

****

DISTRIBUTION TO: T1E1.4 Technical Subcommittee Working Group Members

****

ABSTRACT

Constellation shaping is a technique which reduces the average transmitted power by construction of a spherical boundary region for a block of symbols. Low-complexity techniques have been devised which provide roughly 1 dB of gain. This adds to the coding gain obtained from the use of error-control codes, at the cost of some expansion of the constituent one or two dimensional constellations. This gain is much easier to achieve than additional gains from increased code complexity, and comes with a very low cost in latency. Further, precoding techniques have been developed which preserve the shaping gain. In this contribution we outline some of the alternative shaping and precoding techniques which may be suitable for SHDSL.

Shaping Gain for Single-pair High-speed Digital Subscriber Lines

1. Introduction

To achieve the same reach as two-pair HDSL in the presence of disturbers such as HDSL and ADSL, in SHDSL advanced digital communications techniques will be required, so that the system operates near the Shannon channel capacity. Similar challenges have been faced in the development of the V.34 modem, with the resulting standard being the result of the collective efforts of many researchers in a variety of companies and universities. These techniques represent the state of the art in single-carrier digital transmission, and thus are of interest to the SHDSL community. In this contribution we outline how a combination of channel coding, constellation shaping, and precoding may be of use in providing large performance improvements over uncoded modulation, and discuss some of the research issues that arise for SHDSL in the choice of constellation shaping method and precoder. In Section 2 we describe the concept of shaping, the gains to be expected, and some candidate shaping methods. In Section 3 we discuss precoding, including methods that preserve shaping gain. We conclude with research questions for the SHDSL community.

2. Shaping

Constellation shaping is a technique which reduces the average transmitted power while preserving the minimum distance of coded or uncoded systems. In the methods considered for V.34 this is accomplished by grouping N D-dimensional channel symbols together, and choosing collections of symbols from within an approximately spherical region in ND dimensional space. Because signal points which uniformly fill a sphere have less average power than signal points which uniformly fill a cube, improved average power results. The difference in power between a cube and a sphere grows with dimension, achieving a maximum value of 1.53 dB as the dimensionality tends to infinity. One cost is that the constituent one- or two-dimensional constellations must have more signal points; in the limit as dimensionality goes to infinity, the selection probabilities for the individual signal points follow a Gaussian distribution, with low-energy points selected with higher probability. However, in practice the expansion factor for the constituent constellations can be made small (25 to 50%) while obtaining almost all of the possible gain. The other major cost of shaping is that addressing the points in a spherical region is more difficult than addressing points in a cubic region, since for the latter we may address coordinate by coordinate. This difficulty grows with dimensionality. Thus effective shaping techniques must balance shaping gain, constellation size, and addressing complexity.

There are several different approaches to constellation shaping. In [1] the shaping region is chosen as a scaled version of the Voronoi region of a high-dimensional lattice. Decoding methods for the corresponding lattice code can then be used for addressing the constellation. This amounts to a direct approach to approximating spherical boundaries. An alternative method is proposed in [2], where the constituent constellation is divided into a series of concentric rings of increasing energy. Then a code is used to select a sequence of rings so that the average power is minimized. This approach is in essence a direct implementation of the notion that the selection probabilities for the signal points should follow a truncated Gaussian distribution if a large gain is to be obtained. Constellation expansion is determined by the rate of the code. This work established that 8 to 16 rings are sufficient to achieve almost all of the theoretically possible shaping gain for any given dimensionality of code, and further presented methods to make the addressing complexity for shaping essentially independent of the constellation size.

The technique actually used in the V.34 modem is known as shell mapping, based on a scheme proposed in [3], with related methods discussed in [4-7]. This technique has low latency, a good trade-off between constellation expansion and shaping gain, low complexity, and is easily implemented in combination with the precoding techniques outlined further. Variations on shell mapping thus appear to be the most promising for SHDSL.

Shell mapping is a systematic addressing technique which identifies those points which lie within an approximately spherical region, with the property that under a suitable transformation their coordinates may be independently addressed. This greatly reduces the size of the look-up table required to address the multi-dimensional points. The points are also selected to reduce the constellation expansion ratio as required As implemented in V.34, between 0.8 to 1.0 dB of gain results (depending on constellation size) with only 25% constellation expansion, a look-up table with 234 16-bit entries, and 42 multiply/accumulate operations per two-dimensional symbol, for QAM constellations with as many 960 points [3]. A similar technique reported in [4] for 50% constellation expansion yields 1.15 dB with 45 multiply/accumulate operations and a 1.5 kbyte look-up table.

The basic technique is to divide the constituent constellation into a set of q rings, each containing an equal number of points (8 rings in V.34 [8]). Then the shell mapper selects a sequence of N rings over N symbols, while uncoded bits and the trellis encoder select the actual constellation points within these rings. To each ring is assigned a cost, in V.34 just the integer index of the ring started from the lowest power ring (cost=0) to the highest (cost=7). Alternatively, with somewhat higher complexity but slightly improved gains, a cost reflecting the actual average power of the points in the ring could be used. A shell consists of all those combinations of rings with the same sum of the costs of individual rings. The idea is to preferentially select the low-cost shells, while avoiding a global look-up table listing all the low cost combinations. The techniques reported in [3-7] all are motivated by the need to break down the multidimensional addressing problem into a series of low-dimensional addressing methods, using results both from the vector quantizer literature and techniques for addressing shells of lattices [9]. The result is a mix of computations, some to perform a mapping of the shells onto regions which may be addressed as combinations of low dimensional regions (viewed in a different space), and algorithms to actually perform the (simple) addressing of these constituent parts. The algorithms, while conceptually difficult, are easy to actually implement and well-suited to fixed point arithmetic implementations.

Given the state of maturity of research in this area, it is not likely that further large reductions in complexity will result for new techniques. However, the trade-offs for SHDSL are different from V.34 in a number of respects. First, the number of information bits per symbol is far less than in V.34 (2-3 bits per dimension). Shaping must act on those bits which do not pass through the trellis encoder, or else the coset sequence is modified. Thus, relatively few rings are possible in one or two dimensions, and so to achieve significant gain consideration might be given to acting on four dimensional rings. In this case, relatively short blocks will be required, and thus the addressing complexity may be quite small. Second, in V.34 constellation expansion due to both trellis coding and shaping is limited by the presence of A/D and D/A converters in the transmission path. In SHDSL this is not an issue, and constellation expansion is chiefly limited by considerations of the linearity of the analog components used to couple the modem to the line. Allowing large constellation expansion can result in either lower complexity or larger gain. Finally, for V.34 the shaping algorithm had to work over a large range of constellation sizes, with the larger ones being the most critical in the design. For SHDSL there will be only one constellation size, and it involves relatively few bits per dimension. All of this argues for a careful consideration of the different shell-mapping techniques.

3. Precoding

The use of adaptive equalization is the single most important factor in improving the speed of voice-band modems, and clearly is also critical in subscriber line applications. Using a decision feedback equalizer (DFE) the symbols may be transmitted at close to the Nyquist rate with far less complexity than maximum likelihood sequence estimation (MLSE). Margin is also improved through the use of trellis coded modulation (TCM). This is especially well-suited to V.34 since the constellations used are both large and variable. With TCM, not every bit needs to be coded, and the same code can be used for all constellations, with no modifications to the decoder. All that needs to be changed is the block which performs modulo reduction of the constellation to perform branch metric calculations, and even these changes are quite simple. The standard specifies the possibility of using any of three possible codes (16, 32 and 64 states). Four-dimensional codes are used to reduce vulnerability to non-linearity through lesser constellation expansion ratios. These codes provide up to 5 dB of gain.

Unfortunately, a DFE does not fit together very well with channel codes. It requires zero delay decisions, which of course must be uncoded, and then may generate bursts of errors. Use of an interleaver to break up these errors yields unacceptable delay and still generates worse performance than a combined MLSE decoder would use. Variations on MLSE which have somewhat reduced complexity such as RSSE perform well on modem channels, but are still higher in complexity than the solution actually adopted, namely precoding. In this technique, in effect the feedback section of the DFE is moved to the transmitter, where it operate on error-free symbols. Modulo reduction following the decision then whitens the transmitted sequence; a similar modulo reduction device is required in the receiver following the forward filter and prior to decoding. This technique has performance comparable to MLSE with a fraction of the complexity, and thus some combination of precoding and trellis coding should be used in single-carrier SHDSL systems.

However, the standard Tomlinson-Harashima (TH) precoding technique described in the next subsection does not preserve shaping gain. Thus, over the course of the V.34 standard development a series of progressively better techniques were developed. The one actually adopted yields a complete separation of coding and shaping gain and ISI rejection, with no more complexity than the standard approach, also supporting the use of rotationally invariant codes. We will briefly explain the TH technique to assess its impact on the complexity of the various components of a SHDSL system, and then discuss precoding techniques which permit shaping gain.

Tomlinson-Harashima Precoding

In this section, we consider the interaction of channel coders, variations on decision feedback equalizers and echo cancellers. A DFE is depicted in Figure 1 below. The feedforward filter (FFF) consists of a tapped delay line, spaced either at the symbol rate or fractions of the symbol rate. It is an FIR filter, whose output is sampled at the symbol rate. The feedback filter (FBF) is also an FIR filter, differing in structure from the FFF in only two ways (1) it is fed decisions, reducing the complexity of the multipliers since the number of signal levels is usually small and (2) there is no feedback coefficient corresponding to the current decision. The FBF is always symbol spaced. Its output is subtracted from that of the FFF, and a slicer then makes decisions on the most likely transmitted symbol.

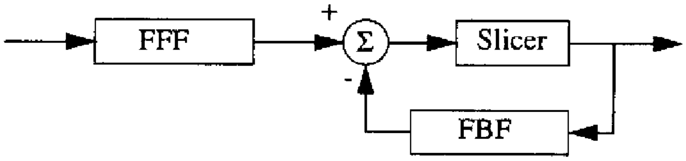

Figure 1: Decision Feedback Equalizer

When the length of the two filters becomes large, the forward filter acts mainly to remove the precursor ISI and whiten the residual noise. In this case, the FFF is just a noise-whitening matched filter. The feedback section then removes all the postcursor ISI (which is typically made more severe by the noise whitening and precursor suppressing operations). When the filters have finite length, the operations of the two filters may not be so neatly separated. The forward filter will shape the noise and residual ISI in such a way that the residual distortion following the action of the FBF is white. Thus, the FFF will to some extent also reduce the postcursor ISI.

The action of the DFE is very close to what a maximum likelihood sequence estimator (MLSE) would do. Its structure is a noise-whitening matched filter (e.g. the FFF of a DFE), followed by a Viterbi decoder, which searches a trellis whose (ISI) states are determined by the coefficients of the FBF of a DFE. The DFE may be viewed as a reduced complexity MLSE structure which can only follow one path through the trellis--the most likely one for a one-step decision. In the absence of decision errors, its performance is close to that of MLSE with a small fraction of the complexity. However, with decision errors it can be far inferior to MLSE, which makes final decisions only after examining a large number of candidate paths.

The use of channel coding does not help the DFE, since the DFE requires zero delay decisions while decoding inherently introduces delay. Trellis coded modulation actually makes matters worse, since the constellation is expanded, reducing the distance between channel symbols. The solution to this problem is to move the FBF to the transmitter, so that only noise-free symbols experience feedback. However, more than this is required, since the use of a simple feedback loop in the transmitter will result in (unpredictable a priori) spectral shaping of the transmitted sequence. In a precoded system as shown in Figure 2, a modulo reduction device is provided within the feedback loop, which brings the transmitted signal level to within the range of the original constellation. This both reduces the peak transmitted power and whitens the spectrum. In practice, the spectrum is nearly indistinguishable from the original modulation. Another modulo device in the receiver following the FFF then serves to recover the equalized signal. The channel decoder can then operate as normal, without concern for error propagation in the DFE. Moreover, because the feedback decisions are all made in the transmitter, there is no possibility of error propagation.

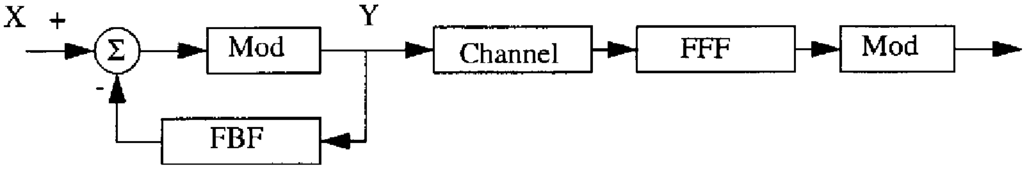

Figure 2: TH Precoded System

There are however several problems that arise with precoding. The first is that the power of the constellation Y is always slightly higher than that of the original constellation X. This penalty is 0.5 dB for a four level constellation X, and rapidly declines as the number of levels grows (less than 0.1 dB for 8 levels). More seriously in an echo cancelled system, the constellation Y takes on a number of levels limited only by the quantization used in the modulo device; in the limit it approaches a continuum. This will dramatically increase the complexity of the multipliers in the echo canceller. Finally, the FBF in the transmitter is more complicated than in the receiver, requiring many more levels in the data values.

The echo canceller and FBF complexity problems may be mitigated by quantizing the output of the modulo device. Then the number of levels is controlled, but this brings on possible problems of error propagation, and issues in training. For example, should the FBF be trained with a quantizer inserted somewhere in the DFE, and if so where? We may proceed as follows. A quantizer following the mod device in the precoder adds a noise sequence to the signal. This is equivalent to adding a quantizer following the FBF in the DFE, and thus to adding a quantizer noise source prior to the slicer. The effect on performance is thus easily analyzed; the quantizer in effect puts a ceiling on the SNR achievable by the precoded system.

Channel decoders typically require soft information of three bits of resolution between the nominal signal points. This produces negligible performance degradation compared to performing real arithmetic. Simulations reveal that quantization to similar levels in the feedback section of the DFE similarly leads to almost no degradation, either in performance or in time to train the equalizer. A simple analysis of the precoder assuming high SNR (zero-forcing) with the quantizer taken to be a noise source produces the same conclusion--the quantizer appears as a noise source at the FBF output, at the same noise level as in the transmitter. Thus, the precoder will demand at worst 3 additional bits of precision in the "data" values for the echo canceller and FBF multipliers.

Precoding and Shaping

One further problem with TH precoders is that they destroy shaping gain by mapping all signals into a rectangular region. Since the additional 1 dB of gain provided by shaping is usually much easier to obtain than a further 1 dB of coding gain, considerable effort went into devising methods that could preserve shaping gain. This first of these was trellis precoding, which combined the operations of trellis coding, trellis shaping, and precoding [10,11]. This technique unfortunately has gain and complexity which depends on the severity of the ISI. Subsequent methods appeared which completely separated ISI performance from shaping gain. However, early versions produced a precoding loss larger than that of TH precoding, limiting the usefulness for small constellations. The version finally adopted ([12], similar to [13]) has the same loss as TH precoding and thus no further improvement can be expected.

The basic principle of operation is as follows. We may view the precoder as adding a dither sequence to the original uncoded sequence. The main conceptual breakthrough comes in realizing that we need only arrange that this dither sequence produce a valid code sequence in the receiver. The dither sequence has a much smaller range than in early forms of precoding which used the Voronoi region of the code (effectively the size of the constellation formed by one representative of each coset) [13]. Rather it is reduced to essentially the Voronoi region of the signal set; reference [12] indicates how to apply the approach to any system using coding based on lattice partitions (e.g. most trellis codes for one-, two- and four-dimensional signals). This approach differs in a further respect from TH precoding: the trellis encoder actually appears in the feedback path of the mapping circuitry. Both uncoded bits and one component of the dither sequence are used in determining the code sequence.

None of these operations has any appreciable effect on system complexity, and it is easy to arrange for rotationally invariant codes to be used. As in TH precoding, the FBF from the DFE appears in the transmitter; in the V.34 standard the degree of quantization for the precoder is specified, limiting the number of levels at the output and thus the echo canceller complexity. Shaping together with constellation expansion from trellis coding will add roughly one bit to the required precision of the operations, depending on the choice of code. Thus, the benefits of the DFE, trellis coding, and shaping may all be obtained with a low complexity solution. The only details that must be changed in the implementation of the precoder for SHDSL compared to V.34 involve the possible use of different lattice partitionings to accommodate different channel codes, and contemplation of much larger feedback filters in the transmitter to deal with the greater channel dispersion.

4. Conclusions

Robust SHDSL transmission at the limits of the CSA will require advanced digital communications techniques such as trellis coding, precoding, and shaping in addition to careful selection of modulations and spectral shaping methods. Techniques which have pushed the performance of voice-band modems towards the Shannon limit are thus very attractive for the SHDSL application. In particular, shaping and precoding methods have been devised that have optimal performance and low complexity, and which work with a large variety of modulations and channel codes. This research can be leveraged for subscriber line applications. What remains to be investigated in detail is which specific combination of these techniques is best suited for SHDSL.

References

[1] G.D. Forney Jr., "Multidimensional constellations II: Voronoi constellations," IEEE J Selected Areas in Comm., pp. 941-958, Aug. 1989.

[2] A.R. Calderbank and L.H. Ozarow, "Nonequiprobable signaling on the Gaussian channel," IEEE Trans. Inform. Theory, vol 36, pp. 726-740, July 1990.

[3] Motorola Information Systems, "Shell mapping for V.fast," CCITT contribution Question 3/XVII WP XVII/1, July 1992.

[4] R. Laroia and N. Farvardin, "On optimal shaping of multidimensional constellations," IEEE Trans. Inform. Theory, vol. IT-40, pp. 1044-1056.

[5] A.K. Kandani and P. Kabal, "Shaping multidimensional signal spaces--Part I: Optimum shaping, shell mapping," IEEE Trans. Inform. Theory, pp. 1794-1808, Nov. 1993.

[6] ---, "Shaping multidimensional signal spaces--Part II: Shell-addressed constellations," IEEE Trans. Inform. Theory, pp. 1809-1819, Nov. 1993.

[7] F.R. Kschischang and S. Pasupathy, "Optimal shaping properties of the truncated polydisc," IEEE Trans. Inform. Theory.

[8] International Telecomm. Union Recommendation ITU-T V.34, "A modem operating at data signalling rates...," Sept. 1994.

[9] R. Laroia and N. Farvardin, "A structured fixed-rate vector quantizer derived from a variable-length scalar quantizer--Part I: Memoryless sources," IEEE Trans. Inform. Theory, May 1993.

[10] G.D. Forney Jr., "Trellis Shaping," IEEE Trans. Inform. Theory, vol. 38, pp. 281-300, Mar. 1992.

[11] M.V. Eyuboglu and G.D. Forney, Jr., "Trellis precoding: Combined coding, precoding, and shaping for intersymbol interference channels," IEEE Trans. Inform. Theory, vol. 38, pp. 301-314, March 1992.

[12] M.V. Eyuboglu, "Device and method for precoding," U.S. Patent 5,446,758, Aug. 1995.

[13] R. Laroia, "ISI Coder--Combined coding and precoding," TIA contribution TR30.2/93-06, Baltimore, June 1993.