

Improved Trellis-Coding for Partial-Response Channels

L. Fredrickson¹, R. Karabed², J. Rae³, P. Siegel², H. Thapar¹, R. Wood¹

¹IBM Storage Systems Division, San Jose, CA 95193

²IBM Research Division, Almaden Research Center, San Jose, CA 95120

³IBM Storage Systems Division, Rochester, MN 55901

Abstract — New code design methods and Viterbi detector architectures are presented for high-rate trellis-coded partial-response (TCPR) systems. The methods, which extend the matched-spectral-null (MSN) coding technique, use novel code constraints and time-varying detector trellis structures to reduce path memory requirements by as much as a factor of two, relative to previously reported codes, while retaining the other attractive features of MSN codes. The design methods and corresponding time-varying trellis structures are illustrated with several examples.

I. INTRODUCTION

Trellis-coded partial-response (TCPR) has been experimentally shown to provide substantial performance advantages over PRML in digital magnetic recording systems. The prototype system evaluated in [1], [2] used a rate 8/10, matched-spectral-null (MSN) code [3] that provided a minimum squared-Euclidean distance 4, a factor of 2 greater than that of the baseline PRML system. Excellent closure between theory, analysis, and the experimental results was achieved. A detailed discussion of the VLSI chip used in the prototype can be found in [4].

From the point of view of hardware implementation, one of the most attractive features of the MSN code used in the prototype was the reduced-complexity, 2-way interleaved, Viterbi detector. On each of the even/odd interleaves, the detector used a time-invariant, six-state trellis structure that tracks the running-digital-sum (RDS) of the code sequences, where the RDS of a binary sequence $\mathbf{a} = a_1, a_2, \dots, a_n$ is defined by

$$RDS(\mathbf{a}) = \sum_{i=1}^n (-1)^{(a_i+1)}.$$

However, this simplified detector trellis supports so-called quasicatastrophic (QC) sequences — infinite sequences that are represented by more than one distinct path through the trellis structure. These sequences give rise to

minimum-distance error events of unbounded length. This, in turn, would degrade the performance of the detector which, in practice, must use finite path memory. To prevent this performance degradation, referred to as QC error-propagation, the rate 8/10 spectral-null code in [1], [2] was further constrained to avoid generating any QC sequences, a property shown in [3] to bound error-event lengths. The maximum length of minimum-distance error events was thereby limited to 42 bits. The length-64 path memory implemented in the prototype system therefore provided ample time for the survivor paths to merge, ensuring that there would be virtually no negative impact on the system performance due to path memory truncation.

Nevertheless, for reasons of hardware complexity, cost, power, and track format efficiency, it remains desirable to further reduce the path memory length required by a TCPR system. Imposing stricter constraints on the code, while retaining the time-invariant detector trellis structure, offers one approach, although this can complicate the code design, increase encoder/decoder complexity, and increase decoder error propagation. As an alternative, this paper introduces evolutionary TCPR design approaches, based upon the use of *time-varying* detector trellis structures, that produce codes suitable for practical magnetic recording systems. (We remark that time-varying structures also appeared in [5] in connection with another trellis-coding approach.)

In Section II, we provide background on MSN codes and elaborate upon the relationship between QC sequences and path memory length requirements. We then give an example to motivate the use of time-varying detector trellis structures to reduce the maximum length of minimum-distance error events. In Section III, we describe several approaches that have been developed to overcome the problems associated with QC sequences, all involving Viterbi detection using time-varying trellis structures. Examples of codes designed with most of these methods are discussed, including a new, rate 8/10 “partitioned” MSN code whose implementation and experimental performance evaluation is the subject of a companion paper [6]. Finally, in Section IV, we offer concluding remarks about the new code design techniques and their use in practical TCPR systems.

Manuscript received August 5, 1994.

L. Fredrickson, now at Quantum Corp., Milpitas, CA 95035; R. Karabed, e-mail karabed@almaden.ibm.com, fax 408-927-3520; P. Siegel, e-mail siegelph@almaden.ibm.com, fax 408-927-3030; J. Rae and H. Thapar, now at DataPath Systems, Inc., Santa Ana, CA 92707, fax 714-546-4245; R. Wood, fax 408-256-2653.

II. BACKGROUND AND MOTIVATION

In the remainder of this paper, we consider trellis codes for the dicode partial response channel, having system polynomial $1 - D$. By interleaving, the codes can be adapted for use, more generally, on channels with system polynomials of the form $1 - D^N$, including the case of most practical interest, the Class-4 partial response channel, where $N = 2$. To illustrate the connection between QC sequences and detector path memory, and to provide motivation for the use of time-varying detector trellises to reduce required path memory length, we will refer to the original rate 8/10 MSN code for the dicode channel disclosed in [7].

A. Quasicatastrophic Sequences and Detector Path Memory

In the context of the binary-input dicode channel, the MSN coding approach is based upon the observation that a dc-free code (i.e., a code whose average power spectrum vanishes at zero frequency) doubles the minimum squared-Euclidean distance relative to the uncoded channel, increasing it from 2 to 4. This translates to an effective coding gain of approximately 3 dB, not accounting for the code rate penalty.

A necessary and sufficient condition for a system of constrained sequences to have a spectral-null at zero frequency is that the sequences have bounded RDS values. Fig. 1 shows a so-called canonical state diagram that generates all binary sequences that assume at most seven RDS values. The states of the diagram effectively track the RDS values of the sequences generated by following paths through the directed graph. Using code construction methods for constrained systems, one can derive from the canonical diagram a finite-state-machine encoder for a rate 8/10 MSN code.

A reduced-complexity trellis structure is derived not from the graph underlying the finite-state-machine encoder, but from the typically much simpler canonical diagram, which we will refer to on occasion as the RDS diagram.

Fig. 2 shows one stage of one component of the 2-step (i.e. 2 symbols per edge) RDS trellis structure corresponding to the RDS diagram of Fig. 1. The state labels correspond to those of Fig. 1, and the 2-symbol label for each edge is indicated to the left of the state from which the edge emanates. The detector trellis incorporates the spectral-null constraint represented in the RDS trellis, as well as the single-bit memory of the dicode channel. Fig. 3 shows the resulting 2-step detector trellis structure. The state labels are of the form s,t , where s is the RDS value and t is the dicode channel memory. For each edge, the channel input-symbol pair u and the corresponding channel output-symbol pair v are indicated as u/v to the left of the state from which the edge emanates.




Figure 1. Canonical diagram with 7 RDS states.




Figure 2. RDS trellis (2-step) derived from Fig. 1.




Figure 3. Detector trellis (2-step) for Fig. 2 on dicode channel.

The RDS trellis in Fig. 2 supports sequences that are generated by multiple, distinct paths, namely, the sequences corresponding to paths which do not reach both state 1 and state 7. We say that these sequences are QC with respect to the RDS trellis. Similarly, the detector trellis in Fig. 3 supports QC channel output sequences, and these correspond precisely to the QC sequences in Fig. 2. For example, the sequence 1 0 1 0 1 0 ..., generated in Fig. 2 by the three paths 1 1 1 ..., 3 3 3 ..., and 5 5 5 ..., produces at the channel output the sequence 1 - 1

1 -1 1 -1 ..., generated in Fig. 3 by the three paths (1,0) (1,0) (1,0) ..., (3,0) (3,0) (3,0) ..., and (5,0) (5,0) (5,0) There are, in fact, uncountably many QC sequences supported by the trellises.

As mentioned above, the presence of QC sequences can have an impact on Viterbi detector performance. At moderate to high signal-to-noise ratios, the probability of a detector error is largely determined by the minimum Euclidean distance between sequences forming an error event, where an error event refers to a pair of channel output sequences generated by paths that diverge from a common state and later remerge at a common state. The theory of MSN coding ensures that the minimum squared-Euclidean distance of an error event in the trellis of Fig. 3 must be at least 4. However, the length of minimum-distance error events is, in fact, unbounded, as illustrated by the events generated by pairs of paths of the form (1,0) (1,0) (1,0) (1,0) ... (1,0) (3,1) and (1,0) (3,1) (3,0) (3,0) ... (3,0) (3,1), with corresponding output sequences 1 -1 1 -1 1 -1 ... 1 0 and 1 0 0 -1 1 -1 ... 0 1. In general, to avoid any performance loss, a Viterbi detector should use a path memory length that exceeds the truncation depth of the trellis, meaning the maximum sequence length required to ensure that any pair of sequences, generated by paths that diverge from a common state, accumulate the minimum squared-Euclidean distance. It is clear from the example just given that in the trellis of Fig. 3 the truncation depth is infinite. It follows that system performance may be degraded if the detector uses a finite-length path memory.

In the case of MSN-coded PR systems, the reduced-complexity trellis often supports sequences that are not generated by encoded channel input sequences. The Viterbi detector must distinguish between a valid encoded channel output sequence and other sequences supported by the trellis, including those not in the range of the coded channel. Therefore, the relevant error events are those "generalized" events in which at least one of the sequences in the pair is an encoded output sequence. Similarly, the path memory length should exceed the "generalized" truncation depth (GTD), whose definition is analogous to that of truncation depth, but with the requirement that at least one of the diverging paths correspond to an encoded channel output sequence.

The rate 8/10 code described in [3] did not avoid all QC sequences in Fig. 2. We now examine this TCPR system and demonstrate that the presence of QC sequences leads to unbounded minimum-distance generalized error-event lengths and infinite GTD.

The underlying structure of the finite-state-machine encoder is represented in Table I. The state designations in the table indicate the states in the RDS trellis (Fig. 2) to which the encoder states correspond. Each row of the table corresponds to an encoder state, as indicated in the first column. Successive columns represent possible next states.

State	3a	3b	5a	5b
3a	—	—	A(128)	A(128)
3b	D(132)	—	B(62)	B(62)
5a	\bar{A} (128)	\bar{A} (128)	—	—
5b	\bar{B} (62)	\bar{B} (62)	D(132)	—

Table I. Encoder structure for rate 8/10 MSN code.

The entry in row i column j is a symbol identifying a list of allowable 10-bit codewords that originate in state i and terminate in state j . The number in parentheses is the number of words in the list. The overbar notation \bar{X} is used to denote the list of codewords obtained by bit-wise complementing the codewords in list X . The lists are characterized as follows. Let S denote the set of 10-bit sequences in the canonical diagram of Fig. 1 that originate from state 3 and terminate in state 5. Then B is the subset of S consisting of those words whose final two bits are 11. The subset A is simply the complement of subset B in S . The cardinalities of these sets are $|S| = 190$, $|A| = 128$, and $|B| = 62$. Set D is a subset of size 132 of the set of 162 10-bit words that originate and terminate at state 3, but never pass through state 6 or state 7 in Fig. 1. Each encoder state generates 256 codewords, and an appropriate one-to-one assignment of the 256 distinct data bytes to the codewords completes the definition of the encoder.

It can be seen directly that the encoder generates QC sequences. For example, the subset of code sequences included in the concatenation of subsets $DADA$ which correspond to a series of length-10 paths that emanate from state 3, proceed to state 5, and eventually return to state 3, all without passing through state 7, are QC and could be generated by a parallel path beginning in state 5.

Now, consider the family of code sequences of the form (1 0 1 0 1 0 1 0 1 0) [(1 0 1 0 1 0 1 0 1 0) (0 0 1 0 1 0 1 0 1 0)] ... [(1 0 1 0 1 0 1 0 1 0) (0 0 1 0 1 0 1 0 1 0)] [(1 0 1 0 1 0 1 0 1 0) (0 0 1 0 1 0 1 0 1 0)] generated by the sequences of encoder states 3b 3a [5a 3a] ... [5a 3a] 5b 5a; and the family of sequences, also supported by Fig. 1, of the form (1 0 1 0 1 0 1 0 1 1) [(1 0 1 0 1 0 1 0 1 1) (0 0 1 0 1 0 1 0 1 1)] ... [(1 0 1 0 1 0 1 0 1 1) (0 0 1 0 1 0 1 0 1 1)] [(1 0 1 0 1 0 1 0 1 1) (0 0 0 1 0 1 0 1 1)] generated by the sequences of length-10 paths with corresponding sequence of originating states 3 5 [7 5] ... [7 5] 7 5. The corresponding minimum-distance generalized error events, diverging at state (3,0) and remerging at state (5,1) in Fig. 3, have unbounded lengths,

implying that the GTD of this code with respect to the detector trellis is infinite.

A careful examination of the encoder reveals that, in fact, the duplicate path representations of all of the QC encoded sequences have the property that they pass through state 1 or state 7 at times 0 (modulo 10). We will use this observation in the next section to illustrate how the substitution of a simple time-varying detector trellis structure produces a system with finite GTD.

B. Motivation for Time-Varying Trellis Structures

As mentioned above, it was shown in [3] that one can overcome the problem of infinite GTD by designing the MSN encoder to avoid the QC sequences supported by the trellis, and, in principle, there need not be a penalty in terms of rate loss. However, as subsequent constructions of rate 8/10 MSN codes have demonstrated, it is quite difficult using the constrained coding approach to reduce the maximum error-event length to 40 code bits or less, without substantially increasing encoder and decoder complexity.

The alternative approach, which we now illustrate, is the use of time-varying trellis structures to facilitate elimination of QC sequences from the TCPR system. From the characterization above of the duplicate trellis path representations for QC sequences in the original, rate 8/10, TCPR system, one can verify that by deleting states and edges from the detector trellis — namely the states (1,0) and (7,1) at times 0 (modulo 10) corresponding to the codeword boundary, along with all edges entering or leaving those states — these code sequences will then be generated by a unique trellis path. In other words, the formerly problematical code sequences are no longer QC with respect to the modified time-varying trellis. Moreover, the modified trellis still supports all of the code sequences. It follows from [3] that the TCPR system using the same code with the time-varying detector structure has only a finite path memory length requirement. The new trellis code design methods described in the next section extend this approach and confirm the usefulness of time-varying detector trellis structures as a practical means for substantially reducing the GTD and path memory required in TCPR systems.

III. TIME-VARYING CODE DESIGN METHODS

In this section, several new code design methods for TCPR will be discussed and illustrated. Each method can be used to generate high-rate codes with reduced truncation depth, and the selection of design approach in any particular instance will depend on the relative importance of path memory length, code runlength constraints for timing and gain control, encoder/decoder complexity, and trellis complexity. In all of these examples, the trellis construction will be obtained by modification of RDS

detector trellis structures that combine the RDS constraints with the memory of the $1 - D$ channel.

A. Pruned RDS Trellis Structures

The first class of codes, and the simplest conceptually, is obtained by pruning edges from the canonical MSN detector trellis for a spectral null at zero frequency, in analogy to the motivating example above. We refer to the resulting trellises as “pruned RDS trellis structures” [8]. Deleting states and edges from the trellis may reduce the achievable code rate and, as we shall see in the examples, there may be other costs associated with designing a high-rate code for the pruned trellis structure, such as larger codeword length or increased decoder error propagation.

Perhaps the simplest examples of this kind are the dc-free, or balanced, block codes, in which every codeword $c = c_1, \dots, c_{10}$ has a net RDS equal to 0, $RDS(c) = 0$. A trellis structure supporting balanced sequences of length $2n$ forces all sequences to a single RDS state every $2n$ bits. This restriction clearly eliminates all QC sequences at the channel output. For a given constraint on the number of allowable RDS states, and a given codeword length, one can determine the maximal number of balanced codewords, and therefore the largest achievable block code rate.

For example, to design a balanced block code with rate 4/5, with the number of RDS states limited to seven as in the previously designed rate 8/10 MSN codes — which had multi-state, as opposed to block, encoders — the smallest codeword length one can use is 20 bits. The pruned RDS trellis structure, which supports 107,616 length-20 balanced words, is shown in Fig. 4. A rate 16/20 code results from selecting $2^{16} = 65,534$ codewords from those generated by the RDS trellis paths. Note that rate 16/20 is also the highest rate for any block code supported by this trellis.




Figure 4. Pruned RDS trellis for rate 16/20 MSN code.

The corresponding detector trellis, obtained by combining the decode channel memory with the constraints in the RDS trellis, will require no more than 6 states (ACS-units) at any stage. The detector trellis structure ensures that the minimum-distance error-event length for

any of the candidate rate 16/20 coded systems will not exceed 20 code bits, confining the impact of any minimum-distance event to two data bytes. One can also verify that a GTD of 23 code bits can be achieved.

This approach may also be used to design block codes that, although not balanced, require each codeword to have a constant RDS. Such a code has a null at zero frequency in the continuous part of the average power spectral density, but a non-zero discrete spectral line at zero frequency. These codes, sometimes referred to as spectral-density null codes, are known to provide the same coding gain as dc-free codes on the dicode channel. Fig. 5 shows an RDS trellis that supports a rate 12/15 spectral-density null code. The truncation depth is 15 code bits.




Figure 5. Time-varying RDS trellis for rate 12/15 spectral-density null code.

We remark that the “pruned RDS trellis” design approach is also effective in the context of coding for quaternary-input partial-response channels. Fig. 6 shows an RDS trellis structure that supports a unique quaternary, rate 8B/6Q MSN code for the dicode channel; that is, the trellis supports exactly 256 codewords emanating from each initial state. (The notation 8B/6Q indicates that 8 binary data symbols are encoded into 6 quaternary code symbols.) The truncation depth for minimum distance is 9 symbols. It is interesting to compare this TCPR system with the rate 8B/6Q codes constructed in [9], where time-invariant trellis structures were assumed. For those codes, the path memory requirement is 24 symbols, and the encoder/decoder functions are considerably more complex than those of the new code. It should be mentioned, however, that this new quaternary code has larger average transmitted signal power than the comparable 8B/6Q codes in [9].

B. Modulo N RDS Trellis Structures

Another design approach is based on the observation that coding gain can be achieved by tracking the RDS values modulo N and constraining the allowable value periodically. Specifically, as described in [10], a conventional RDS trellis is first replicated N times. The N copies of the trellis are interconnected in such a way that

each state is associated with a unique value of the RDS modulo N .




Figure 6. Time-varying RDS trellis for quaternary rate 8B/6Q MSN code.

The particular value of the modulus N is chosen so as to ensure that, when channel inputs are runlength constrained, divergent sequences which constitute a minimum-distance generalized error-event lead to distinct states in the RDS modulo N structure. The RDS modulo N trellis is then pruned to enforce a constraint on the allowable values of RDS modulo N periodically, for some period n , thereby eliminating QC sequences and ensuring increased minimum distance. This approach can be used to design high-rate codes of moderate complexity with attractive runlength constraints. The resulting codes, however, will no longer generate a spectral null, and the pruning of edges must typically be done in an *ad hoc* manner. For further details and examples, the reader is referred to [10].

C. Permuted State Structures

The next method introduces the notion of survivor metric permutation and reassignment. Sequences of a specified length with prescribed RDS or RDS modulo N values are concatenated according to rules based upon the RDS value of a sequence and its successor. The detector trellis structure is derived by combining subtrellises that track the pertinent value and reassigning survivor metrics and survivor sequences, according to the sequence concatenation rules. This reassignment amounts to a periodic permutation of state survivor metrics and sequences, from which the class of “permuted trellis codes” derives its name [11]. The state permutation in the detector trellis is chosen, in a somewhat *ad hoc* manner, to ensure that channel memory is preserved, coding gain is achieved, QC sequences are eliminated, and acceptable runlength constraints are imposed. However, the sequences generated by this permutation of trellis states will not have a spectral null, in general.

As an example, Fig. 7 shows a time-varying RDS subtrellis of length 3, representing sequences that, begin-

ning in RDS state 5, accumulate RDS +1 or -3, or, beginning in RDS state 3, accumulate RDS -1 or +3.




Figure 7. RDS trellis for rate 2/3 permuted trellis code.

The detector trellis for the corresponding outputs of the decode channel, represented in systolic form [12], is shown in Fig. 8. Trellis states with channel memory 0 are represented by a white square, while those with channel memory 1 are represented by a shaded square. The sequence concatenation rule, indicated schematically in the figure, is specified by requiring that sequences ending in RDS state 6 (respectively, 2) must be followed by a sequence generated from RDS state 5 (respectively, 3), while preserving the channel memory.

In Fig. 8, the sequences corresponding to trellis paths diverging from states B0, C0, D0, E0, B1, or E1 accumulate distance at least 2 by the end of the trellis. The paths diverging from states C1 (or D1), and ending at states B3 and E3, however, generate output sequences that accumulate only distance 1. Since channel memory is to be preserved in the state reassignments, the possible state pairs in the subsequent adjacent subtrellis to be mated with {B3,E3} are {C0,B0}, {C0,D0}, {E0,B0}, and {E0,D0}. Of these, all but {C0,D0} provide path extensions that remerge after accumulating additional distance of only 1.

The pair {C0,D0}, however, ensures that an additional distance of at least 3 is accumulated by all possible path extensions. If channel memory is to be preserved, this assignment then forces the state pair {A3,F3} to be mated with {B0,E0}. The resulting trellis structure with state permutations has minimum error-event distance equal to 4. Additionally, the trellis structure resulting from the chosen state permutation supports no QC sequences. This follows from the fact that 0 0 0 is the only output sequence generated by distinct paths in the subtrellis (namely, path B0, C1, D2, E3 and path E0, D1, C2, B3)

and the state permutations prevent the repetition of this sequence.




Figure 8. Detector trellis for rate 2/3 permuted code.

Table II describes a 2-state, rate 2/3 encoder that generates sequences supported by the permuted RDS trellis. The path memory required for this system is no more than 12 code bits.

State	Encoder Input	Encoder Output	Next State
X	00	100	X
X	01	001	X
X	10	010	X
X	11	111	Y
Y	00	011	Y
Y	01	110	Y
Y	10	101	Y
Y	11	000	X

Table II. Encoder for rate 2/3 permuted code.

The permuted code approach has been used to design high-rate trellis codes for the decode channel. Fig. 9 shows a detector trellis for a rate 8/10 permuted code with minimum squared-Euclidean distance 4 on the decode channel. The trellis supports exactly the output sequences generated by a 2-state rate 8/10 encoder (not shown). State reassignment occurs every ten bits, at the codeword boundary. Rate 8/10 permuted trellis codes with maximum length of minimum-distance generalized error-events no more than 20 code bits have been designed.




Figure 9. Detector trellis for rate 8/10 permuted code.

It should also be noted that state permutation may be invoked within a codeword. This added flexibility can be useful in designing codes with improved runlength constraints.

D. Partitioned Trellis Structures

Unlike the modulo N trellis codes and permuted trellis codes, the next class of codes — denoted “partitioned MSN codes” — achieves comparable or better truncation depth without sacrificing the spectral null property [13]

As in the design of permuted codes, the candidate codewords of a partitioned MSN code are defined by certain concatenation rules applied to RDS subsequences. The rules are selected to maintain a spectral null while, for any given initial value, constraining the RDS values achieved at the concatenation points. If one views the sequences as lying on an RDS trellis pruned to constrain the RDS states at the beginning and end of codewords to a specified subset, the effect of the concatenation rules is to “partition” the QC sequences among the initial states. By making these sequences available for use as codewords, the encoder design for a specified code rate is simplified.

For example, Fig. 10 shows a concatenation rule for subsequences of length 5. The subsequences emanating from initial state 3 are those sequences with RDS 1 or RDS 3 (i.e., those ending in states 4 or 6), while those emanating from initial state 5 are those with RDS -1 or RDS -3 (i.e., those ending in states 4 or 2).

This concatenation scheme produces 275 candidate 10-bit codewords from each of the two initial states, from which 256 may be selected to create a two-state encoder for a rate 8/10 MSN code. This code is invariant to 180 degree phase-shift; the maximum runlength of zero symbols at the decode channel output is 5; and the encoder has a block decoder (i.e., no error propagation).




Figure 10. Subsequence concatenation rule for rate 8/10, partitioned MSN code.

The partitioning of QC sequences is reflected in a trellis structure obtained by a corresponding “partitioning” of the pruned RDS trellis structure, as shown in Fig. 11.




Figure 11. Partitioned detector trellis structure.

The resulting rate 8/10, partitioned MSN TCPR system has GTD equal to 20 code bits, the best of any of the previous systems. Moreover, the maximum length required for diverging pairs of trellis sequences, in which at least one is a code sequence, to achieve the next-minimum distance (6) is only 30 bits, superior to any other known rate 8/10 code.

This partitioned structure certainly increases detector complexity relative to the pruned RDS trellis: the maximum number of ACS units required has been increased to 10, and certain ACS units at the concatenation point require 4-way compare-select operations. Nevertheless, there is an architecture for the Viterbi detector that may

be used to simplify the implementation of the detection algorithm, as described in the companion paper [6].

IV. CONCLUDING REMARKS

We have presented several new design methods for trellis-coded partial-response (TCPR) systems. The methods were illustrated by examples that demonstrate the opportunities to balance various system parameters such as encoded runlength and spectral-null constraints, encoder/decoder complexity, detector trellis complexity, and path memory length. In particular, all of the methods permit the design of rate 4/5 codes that significantly reduce path memory requirements, relative to previously published rate 8/10 MSN codes, while maintaining other desirable features along with reasonable implementation complexity. The proposed codes and time-varying detector trellis structures should enable channel designers to address all of the practical issues related to the deployment of TCPR in magnetic recording systems, and to achieve in a disk drive the many performance benefits that have been consistently and unequivocally demonstrated in two generations of laboratory prototypes.

REFERENCES

- [1] H. Thapar, J. Rae, B. Shung, R. Karabed, and P. Siegel, "On the performance of a rate 8/10 matched spectral null code for Class-4 partial response," *IEEE Trans. Magn.*, vol. 28, no. 5, pp. 2884-2889, September 1992.
- [2] H. Thapar, B. Shung, J. Rae, R. Karabed, and P. Siegel, "Real-time recording results for a trellis-coded partial-response (TCPR) system," *IEEE Trans. Magn.*, vol. 29, no. 6, pp. 4009-4011, November 1993.
- [3] R. Karabed and P. Siegel, "Matched spectral null codes for partial response channels," *IEEE Trans. Info. Th.*, vol. 37, no. 3, pt. II, pp. 818-855, March 1991.
- [4] B. Shung, P. Siegel, H. Thapar, and R. Karabed, "A 30-MHz trellis codec chip for partial-response channels," *IEEE J. Solid-State Circuits, Special Issue on Analog and Signal Processing Circuits*, vol. 26, no. 12, pp. 1981-1987, December 1991.
- [5] A.R. Calderbank, C. Heegard, and T.A. Lee, "Binary convolutional codes with application to magnetic recording," *IEEE Trans. Info. Th.*, vol. IT-32, no. 6, pp. 797-815, November 1986.
- [6] J. Rae, G. Christiansen, H. Thapar, S. Shih, P. Siegel, and R. Karabed, "Design and performance of a VLSI 120 Mb/s trellis-coded partial response channel," *Proceedings of 1994 The Magnetic Recording Conference*, San Diego, California, August 1994.
- [7] R. Karabed and P. Siegel, "Matched Spectral Null Trellis Codes for Partial Response Channels," U.S. Patent No. 4,888,779, issued December 19, 1989.
- [8] L. Fredrickson, R. Karabed, J. Rae, P. Siegel, H. Thapar, and R. Wood, "Time-Varying Viterbi Detector for Control of Error Event Length," U.S. Patent No. 5,280,489, issued January 18, 1994.
- [9] E. Eleftheriou and R. Cideciyan, "On codes satisfying Mth-order running digital sum constraints," *IEEE Trans. Info. Th.*, vol. 37, no. 5, pp. 1294-1313, September 1991.
- [10] L. Fredrickson, "Time-Varying Modulo N Trellis Codes for Input Restricted Partial Response Channels," U.S. Patent No. 5,257,272, issued October 26, 1993.
- [11] L. Fredrickson, R. Karabed, P. Siegel, and H. Thapar, "Permuted Trellis Codes for Input Restricted Partial Response Channels," U.S. Application Serial No. 08/174,904, filed December 24, 1993.
- [12] L. Fredrickson and J. Rae, "Systolic Array Architecture for Trellis-Code Viterbi Detectors," U.S. Patent No. 5,327,440.
- [13] R. Karabed, J. Rae, P. Siegel, and H. Thapar, "Matched Spectral Null Codes With Partitioned Systolic Trellis Structures," U.S. patent application filed, August 1994.