9600651-5 2005

REPLACED SHEETS SG
☑ Description : paged ...1...1... to paged ...10...
☑ Claim(s) : paged ...11... to paged ...13...
☑ Drawing(s) : paged ...14... to paged ...17...
☑ Abstract : paged ...18... to paged ...-...
DATE OF RECEIPT .....

MULTILEVEL CODING USING TRELLIS-CODED MODULATION AND REED-SOLOMON CODES

Background of the Invention

The present invention relates to multilevel coded modulation useful, for example, in voiceband data transmission (e.g., modem) applications.

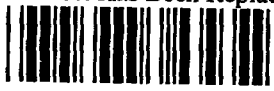
As used herein, the term "multilevel coded modulation" refers to arrangements in which input data is divided into two or more streams which are individually encoded using respective redundancy codes. The coded outputs are then used jointly to select channel symbols from a predetermined signal constellation for transmission over a communication channel, such as a voiceband telephone channel. The principal advantage of adopting a multilevel coded modulation approach is that it provides the system designer with increased flexibility in designing a coding scheme which provides desired levels of error-rate performance, or "coding gain," while meeting various constraints on code complexity and decoding delay.

Illustrative articles describing multilevel codes are: A. Ushirokawa et al, "Multilevel Codes for High-Speed Voiceband Data Modem," Proc. of IEEE Globecom 1989, pp. 1971-5; J. Wu et al, "Multi-Level Multidimensional Trellis Codes," International Symposium of Information Theory, Abstracts of Papers, sponsored by IEEE Information Theory Society, 1990, p. 110; and Pottie et al, "Multilevel Codes Based on Partitioning," IEEE Transactions on Information Theory Vol. 35, No. 1, January, 1989, pp. 87-98. The Pottie article, in particular, suggests in a very general way that the redundancy codes--referred to in the article as the "component codes"--can be selected from any of a number of known types, including 1) binary block and convolutional codes, 2) coset codes (including Ungerboeck and lattice codes), 3) codes on GF(2^m) (e.g., Reed-Solomon codes), 4) uncoded bits, 5) any single- or multilevel partition code, or 6) the concatenation of class 3 with any other class.

Summary of the Invention

In accordance with the present invention, I have discovered that a multilevel coded modulation scheme using a particular combination of two of these types of redundancy codes can be particularly advantageous. In particular, a first portion of the input data referred to herein as the "trellis-encoded bits" is trellis encoded, and the resulting encoded stream is used to identify for each of a succession of symbol intervals a particular one of a predetermined number of subsets of symbols of a predetermined constellation. The remaining input data referred to

This Sheet Has Been Replaced



herein as the "non-trellis-encoded bits" is coded using a Reed-Solomon (RS) code whose output is used to select for transmission a particular symbol from the identified subset. Advantageously, large coding gain can be obtained even when the RS code is a high-rate (i.e., low redundancy, highly-bandwidth-efficient) code.

Preferred embodiments use a 2N-dimensional constellation (N = 1, 2, 4, 8. . .) based on a rectangular lattice. The constellation is partitioned into 4N subsets. The number of states of the trellis code is selected based on the desired error-rate performance for the trellis-encoded bits. The error-rate performance of the non-trellis-encoded bits, were they not to be encoded at all, would typically be worse than that for the trellis-encoded bits, so that the error-rate performance of the coding scheme as a whole, i.e., the overall coding gain, would be determined, or "dominated," by the error-rate performance of the non-trellis-encoded bits. The code implemented by the RS coder, however, is such as to increase the error-rate performance for the non-trellis-encoded bits to a level which is at least as great as that provided for the trellis-encoded bits, so that the error-rate performance for the coding scheme as a whole is higher, it being as great as that which the trellis code provides for the trellis-encoded bits. Advantageously, it turns out that very simple high-rate RS codes can be used in order to achieve this.

I have found two particular multilevel coded modulation schemes to be particularly advantageous. The first uses four-dimensional (4D) constellations, partitioned into eight 4D subsets; a 4D, 64-state trellis code; and a single-error-correcting RS code. The second uses two-dimensional (2D) constellations partitioned into four 2D subsets; a 2D, 64-state trellis code; and a double-error-correcting RS code. Either of these schemes may be used to particular advantage in two applications of current interest. One is the transmission of data at about 19.2 Kbps over dial-up telephone channels. The other is the transmission of data at speeds up to the so-called T1 rate of 1.544 Mbps over the so-called telephone subscriber local loop such as is described generally in U.S. Patent 4,924,492 issued May 9, 1990 to R. D. Gitlin et al.

The invention provides quite a number of advantages. For example, the RS code requires very little redundancy to achieve a desired level of overall error-rate performance, or coding gain; the bandwidth efficiency (bits per symbol) can be made to approach that of a system in which the non-trellis-encoded bits are not coded at all; errors due to impulse noise of any magnitude can be corrected; the transmission of data at various different bit rates is easily accommodated; and the multilevel coded symbols selected for transmission can be interleaved in order to

provide enhanced immunity to channel impulse noise and/or correlated noise.

Brief Description of the Drawing

FIG. 1 is a block diagram of the transmitter portion of a telephone voiceband modem utilizing a multilevel coded modulation scheme embodying the principles of the invention;

FIG. 2 is a block diagram of the receiver portion of a telephone voiceband modem capable of receiving and processing the data signals generated by the transmitter of FIG. 1;

FIG. 3 is a chart helpful in understanding some conventional terminology and concepts;

FIG. 4 shows a two-dimensional constellation that may be used in the transmitter of FIG. 1 either by itself or as a constituent of a higher dimensionality, e.g., four-dimensional constellation;

FIG. 5 shows how a four-dimensional constellation used in the illustrative embodiment is partitioned into eight subsets;

FIG. 6 is a table comparing various coded modulation schemes, including multilevel coded modulation schemes embodying the principles of the present invention;

FIG. 7 shows the frame organization used for the Reed-Solomon encoder of the transmitter of FIG. 1;

FIG. 8 shows circuitry which implements a particular trellis code used by the trellis encoder of the transmitter of FIG. 1.

Detailed Description

FIG. 1 is a block diagram of the transmitter portion of a telephone voiceband modem utilizing a multilevel coded modulation scheme embodying the principles of the invention. In overall view, binary data from a data source 101--such as a personal computer--is caused to be represented by 2N-dimensional symbols taken from a predetermined 2N-dimensional signal constellation, which symbols are modulated onto a carrier for transmission over a voiceband telephone channel 150.

Attention is directed briefly to FIG. 3, which will be helpful in understanding some of the terminology and concepts that are conventionally used in this art. Each of the aforementioned symbols is comprised of the concatenation of N constituent two-dimensional (2D) "signal points," N = 1, 2, 3, ... Each such signal point is a point in a predetermined 2D constellation--illustratively shown in FIG. 3 as a so-called QAM constellation. (The number of signal points in the

2D constellation depends on the needs of the application.) A 2N-dimensional symbol is delivered to the transmission channel during N "signaling intervals" of duration T, one signal point in each signaling interval. The assemblage of all the different 2N-dimensional symbols used in any particular coded modulation scheme is referred to as the "2N-dimensional constellation."

In the illustrative embodiment of FIG. 1, the value of N is 2. That is, the signal constellation is a four-dimensional (4D) constellation comprised of symbols taken from first and second 2D signal constellations in the first and second signaling intervals of each 4D symbol interval, respectively. Illustratively, the same 2D constellation is used for both signaling intervals. That 2D constellation, in particular, is illustratively the 64-signal-point (64-point) QAM constellation shown in FIG. 4. Additionally, all possible combinations of two 2D signal points are used in this embodiment, so that the 4D constellation is comprised of $64^2 = 4096$ 4D symbols.

Returning now to FIG. 1, the stream of bits from source 101 is clocked into scrambler 104 at an average rate of 10.8875 bits per 4D symbol interval. (The significance of this rate will be made clear hereinbelow.) Scrambler 104 randomizes the data in conventional fashion. The serial bit stream output of scrambler 104 is applied to serial-to-parallel (S/P) converter 105, which provides 11-bit output words on leads 108/109 for each 4D symbol interval. (As will be clear from the context, various ones of the leads shown and described herein, such as lead 108 or lead 109, will be understood as being, in actuality, a bundle of leads, each of which carries a respective bit.) In particular, two of the bits are provided on lead 108 and the other nine are provided on lead 109. As will be described in detail hereinbelow, S/P converter 105 occasionally will provide only the two bits on lead 108 without providing any bits on lead 109.

The stream of bits on leads 108/109 are applied to an encoder 11 comprised of 4D, 64-state trellis encoder 112 and rate-158/160 Reed-Solomon (hereinafter RS) encoder 114. In particular, the stream of bits, comprising a succession of bit pairs, on lead 108 are supplied to trellis encoder 112, whose output on lead 113 comprises three bits. These three bits identify one of eight predetermined subsets of the 4096 4D symbols of the 4D constellation. The symbols are assigned to subsets in the following, standard way: Each of the two 2D constituent constellations (FIG. 4) of the overall 4D constellation is partitioned into four 2D subsets--denoted a, b, c, and d. FIG. 3 shows by a reference letter which of the four 2D subsets each of the 2D points is assigned to. The eight subsets

of the overall 4D constellation are then arrived at as shown in FIG. 5. In particular, 4D subset S0 is comprised of each 4D symbol in which the first and second constituent 2D signal points are both taken from either 2D subset a or 2D subset b. These combinations of signal points are denoted in FIG. 5 by (a,a) and (b,b), each of which is referred to as a "4D type." Each of the other 4D subsets, S1 through S7, is also formed by combining 2D subsets, as indicated in the FIG. Thus, as another example, 4D subset S3 is comprised of each 4D symbol in which the first and second constituent 2D signal points are taken from 2D subsets a and d, respectively--the 4D type labeled (a,d)--or from 2D subsets b and c, respectively--the 4D type labeled (b,c). Since there are 4096 4D symbols overall and eight subsets, each 4D subset contains 512 4D symbols.

In prior art, conventional trellis-coded modulation (TCM) schemes, the bits provided on lead 109 are so-called "uncoded" bits which are used to select for transmission a particular symbol from the 4D subset identified by the bits on lead 113. Thus in a conventional TCM scheme, each nine-bit word on lead 109 would be used to select one of the $2^9 = 512$ 4D symbols of the identified 4D subset.

In accordance with the invention, however, the stream of bits on lead 109 is not used to select a symbol directly. Rather, those bits are first encoded by RS encoder 114, and it is the output of RS encoder 114--which is still in the form of nine-bit words--that is used to select a particular symbol from the identified subset. The overall coding scheme, then, is a "multilevel coded modulation" scheme in that there are multiple--in this case, two--levels of input bits being encoded. Specifically, some of the bits are trellis-encoded and the rest--the so-called "non-trellis-encoded" bits--are Reed-Solomon encoded.

RS encoder 114 is of a known type--illustratively, a conventional, rate-k / k+2 systematic encoder over $GF(2^9)$ with k = 158. Reed Solomon coding and decoding is described, for example, in Michelson et al, Error Control Techniques for Digital Communication, Chapter 6, John Wiley and Sons, 1985. As such, encoder 114 provides its outputs in RS frames. As shown in FIG. 7, each RS frame is comprised of 160 (i.e., k+2) nine-bit RS words on lead 115. Since the RS code is a so-called systematic code, the first k = 158 words of the frame are simply 158 successive input words from lead 109. These are referred to herein as the "information-bearing words." The last two words of the 160-word frame are so-called "redundant words" generated in response to the values of the first 158 words in accordance with the selected RS code. When the overall frame of 160 words is first recovered in the receiver, the presence of these two redundant words therein

makes possible the identification and correction of any single, erroneously recovered one or two erased ones of the 160 words. This particular RS code is thus referred to as a single-error-correcting RS code. The operation of RS encoder 114 is synchronized with that of S/P converter 105 in such a way that a nine-bit word is provided on lead 109 for each of the first 158 successive 4D symbol intervals comprising a frame and no bits are provided in the remaining two 4D symbol intervals. It is during these two intervals that encoder 114 outputs the aforementioned two redundant words.

As previously described, three bits are supplied on lead 113 for each of the 160 4D symbol intervals. Twelve bits are thus supplied on leads 113 and 115 for each 4D symbol interval--three of the bits (lead 113) identifying a 4D subset and the remaining nine of the bits (lead 115) selecting a particular symbol of that subset. Those bits are provided to 4D, 64-QAM constellation mapper 120, which outputs representations (e.g., the x and y coordinates) of the two constituent 2D signal points of the selected 4D symbol. Those representations are applied to conventional modulator 141, which applies to channel 150 a passband data signal representing those 2D signal points.

It can now be seen why it is that data source 101 is clocked so as to supply its data at the average rate of 10.8875 bits per 4D symbol interval, as noted above. Of the twelve bits needed to select a particular one of the 4096 4D symbols of the constellation, one redundant bit is introduced by the trellis encoder and an average of 0.1125 (= 9 bits x 2/160) bits, i.e., the bits of the two redundant words, are introduced by the RS encoder. As a result, the data rate for source 101 needs to be (12 - 1 - 0.1125) = 10.8875 bits per 4D symbol interval.

We turn, now, to the receiver of FIG. 2.

The receiver receives from channel 150 the passband data signal generated by the transmitter of FIG. 1. The signal is first applied to equalizer/demodulator circuitry 210 which, in conventional fashion, recovers a sequence of signal points which it provides on lead 211 to decoder 22 and, more particularly, to maximum-likelihood decoder 220 therein. Because of distortion and other channel anomalies that circuitry 210 is not able to fully compensate for, the signal points on lead 211 are somewhat displaced in 2D signal space from the 2D signal points that were transmitted. As its name implies, the function of maximum-likelihood decoder 220 is a) to determine--based on a knowledge of the trellis code used by trellis encoder 112--what the most likely sequence of transmitted 4D symbols actually was, and b) to provide on leads 221 and 222 eleven bits

corresponding to those 4D symbols, i.e., corresponding respectively to the bits on leads 108 and 115 in the transmitter.

The remainder of the processing performed in the receiver of FIG. 2 is the inverse of processing performed in the transmitter. Thus, in particular, RS decoder 230 within decoder 22 operates on each received frame of 160 nine-bit words on lead 222 to recover the 158 information-bearing nine-bit words therein. In particular, as noted above, the decoder is capable of identifying and correcting any error-corrupted single nine-bit word or two erased words provided by maximum-likelihood decoder 220. The stream of 158 corrected information-bearing words is supplied by RS decoder 230 on lead 232. The eleven bits on leads 221 and 232 are thereafter converted to serial form by parallel-to-serial converter 270, descrambled by descrambler 280, and applied to a data sink 290 which may be, for example, a mainframe computer.

The advantages of the invention can be understood by considering the following:

A given trellis code, when used in a unilevel coded modulation scheme, i.e., one in which the non-trellis-encoded bits are not encoded at all, provides particular respective levels of error-rate performance for the trellis-encoded and non-trellis-encoded bits. For many trellis codes of interest, the error-rate performance of the non-trellis-encoded bits is worse than that for the trellis-encoded bits. Thus, the error-rate performance of the coding scheme as a whole, i.e., the overall coding gain, is determined, or "dominated," by the error-rate performance of the non-trellis-encoded bits. More particularly, the error-rate performance for the non-trellis-encoded bits is here a function of the minimum Euclidean distance between the symbols within a subset.

(As will be well appreciated by those skilled in the art, the error-rate performance of the trellis-encoded bits is principally determined by the minimum Euclidean distance between different valid sequences of 4D symbols selected respectively from different valid sequences of 4D subsets, that distance being referred to herein as the "trellis distance." (A "valid" sequence is one that is allowed by the trellis code to actually occur.) By contrast, the error-rate performance of the non-trellis-encoded bits is principally determined by the smaller of a) the trellis distance, or b) the minimum Euclidean distance between different valid sequences of 4D symbols selected respectively from the same valid sequence of 4D subsets, referred to herein as the "non-trellis distance.")

If the overall error-rate performance is not adequate or acceptable for a given application, the prior art approach is to partition the constellation into a greater number of subsets. Since there are then fewer symbols per subset, the distance between them is thereby increased, the increase typically being such that the overall error-rate performance becomes dominated by the error-rate performance of the trellis-encoded bits. The overall error-rate performance is thereby increased. There is a big price to be paid for this, however. The increase in the number of subsets requires the use of a trellis encoder that has an increased number of trellis states in the finite-state machine that implements the code. The complexity of the maximum-likelihood decoder is roughly proportional to the product of the number of subsets and the number of trellis states. Thus, the complexity of the maximum-likelihood decoder is increased dramatically. Indeed, it may be increased to the point that a practical and/or cost-acceptable implementation of the coding scheme may not be possible.

The present invention, by contrast, uses the above-described Reed-Solomon encoding, rather than an increase in the number of subsets, as the mechanism for increasing the error-rate performance of the non-trellis-encoded bits. Indeed, low-complexity RS codes are available that can increase the non-trellis distance to a level which is greater than the trellis distance, with the result that the overall error-rate performance becomes dominated by the error-rate performance provided by the trellis code to the trellis-encoded bits. The latter will then dominate, thereby providing an increased level of error-rate performance for the coding scheme as a whole. Assuming, in the first instance, that the same trellis code is used, the complexity of the maximum-likelihood decoder, advantageously, stays the same. Moreover, a single-error-correcting RS code, such as is used in the above-described illustrative embodiment, is highly bandwidth efficient. In the above embodiment, for example, the "overhead" caused by the RS code is only 0.1125 bits per 4D symbol, as previously seen.

Since the error-rate performance of the trellis-encoded bits now dominates, one can, at this point, achieve an even greater level of overall error-rate performance, if desired, by using a trellis code which, although having more states, does not require an increased number of subsets. The overall effect, then, is to raise the overall error-rate performance while incurring a relatively modest increase in maximum-likelihood decoder complexity.

The foregoing is illustrated graphically in the table of FIG. 6. We take as a "baseline" the modulation scheme III shown in the table. This scheme is based on the 4096-symbol, 4D constellation described above and partitioned into eight subsets as also described above. As shown in the table, this scheme uses a) a 4D, 16-state trellis code for the trellis-encoded bits--illustratively the one disclosed in L. Wei, "Trellis Coded Modulation with Multidimensional Constellations," IEEE Transactions on Information Theory, July, 1987, and b) no coding for the remaining, non-trellis-encoded bits. For this scheme, the non-trellis-encoded bits dominate the error-rate performance.

The next entry up, modulation scheme II, embodies the principles of the present invention. It uses the same constellation, the same partitioning and the same trellis code as scheme III. Now, however, the non-trellis-encoded bits are encoded using a single-error-correcting RS code, such as the rate-158/160 RS code described above. The table indicates, as discussed above, that the dominant error factor is the trellis-encoded bits. As symbolically represented by the arrow at the left of the table, the overall error-rate performance of this scheme is better than that of scheme III.

The top entry of the table, modulation scheme I, also embodies the principles of the invention, but now uses a more powerful trellis code. Specifically, that code is a 4D, 64-state trellis code which, advantageously, is based on the same eight-subset partitioning of the same 4D constellation. The overall error-rate performance is again improved by virtue of the increased error-rate performance which this code provides for the trellis-encoded bits. Note that, in this case, the trellis distance is still not as great as that of the non-trellis distance, so that the error-rate performance of the trellis-encoded bits still dominates.

The trellis code used by trellis encoder 112 for scheme I is implemented by the circuitry of FIG. 8. In that FIG., each of the elements labeled "2T" is a delay element which delays its input by 2T seconds, i.e., the duration of a 4D symbol interval. The elements labeled "+" are exclusive-OR gates. This is a "systematic" code, so that two of the three output bits on lead 113 are simply the two input bits from lead 108. The third output bit is generated by the circuitry shown in response to the two input bits and delayed versions thereof. As noted above, the three output bits on lead 113 identify one of the eight 4D subsets S0, S1, . . . S7. In particular, the decimal value of the binary word formed by the three output bits--with the upper, middle and lowest bits being the most-, second-most and least- significant bits--defines the subset. For example, the bit pattern 100 appearing on leads 113 identifies subset S4, since binary 100 is equivalent to decimal 4.

The foregoing merely illustrates the principles of the present invention. For example, although particular bit rates, codes, constellations and partitionings are shown herein and particular parameter values (e.g., for N, k, etc.) are used, these are all merely illustrative. For example, various different source bit rates can be accommodated via the use of a) RS words of different size (number of bits) in combination with b) various different constellation sizes, without, in general, the need to change, for example, the encoding/decoding algorithms of the RS and trellis codes nor the RS frame length. Moreover, it may be noted that a 2N-dimensional constellation may be comprised of constituent constellations of any lower order, e.g., a 4D constellation may be comprised of four 1D constellations. Also, certain 2N-dimensional constellations may be such as to not be resolvable into constituent, lower-dimensional constellations. Moreover, although the invention is illustrated herein in the context of a passband transmission arrangement, it can also be used in baseband arrangements.

Moreover, although particular applications in which the coded modulation scheme of the present invention are disclosed, the invention is usable in other applications as well.

Additionally, it may be noted that, although the invention is illustrated herein as being implemented with discrete functional building blocks, e.g., encoders, mappers, etc., the functions of any one or more of those building blocks can be carried out using one or more appropriately programmed processors, digital signal processing (DSP) chips, etc. Thus although each of the various "means" recited in the claims hereof may correspond, in some embodiments, to specific circuitry which is specifically designed to perform the function of just that means, it will be appreciated that such "means" may alternatively correspond, in other embodiments, to the combination of processor-based circuitry with stored program instructions which cause that circuitry to perform the function in question.

It will thus be appreciated that those skilled in the art will be able to devise numerous and various alternative arrangements which, although not explicitly shown or described herein, embody the principles of the invention and are within its spirit and scope.

Claims:

1. Transmitter apparatus comprising

first and second redundancy encoders (112, 114),

means (101, 104, 105) for applying first and second streams of data to said first and second redundancy encoders, respectively, and

mapping means (120) for providing channel symbols selected from a predetermined signal constellation in response to the outputs of said encoders,

CHARACTERIZED in that said first and second encoders are a trellis encoder and a Reed-Solomon encoder, respectively.

2. The invention of claim 1 wherein said signal constellation is comprised of a plurality of subsets of channel symbols, wherein said mapping means a) identifies a sequence of said subsets in response to the output of said trellis encoder, and b) selects a channel symbol from each subset of said sequence in response to the output of said Reed-Solomon encoder.

3. The invention of claim 2 further comprising means (141) for generating a passband signal representing the channel symbols provided by said mapping means and for applying said passband signal to a transmission channel.

4. The invention of claim 2 wherein said mapping means and said trellis encoder, in combination, provide a first particular level of error rate performance for said first stream of data and wherein said mapping means, said trellis encoder and said Reed-Solomon encoder, in combination, provide a second particular level of error-rate performance for said second stream of data which is at least as great as said first particular level of error rate performance.

5. Receiver apparatus for processing a received signal that was generated by applying first and second data streams to a trellis encoder (112) and a Reed-Solomon encoder (114), respectively, to provide first and second encoder output streams, respectively, and by thereafter providing, in said generated signal, channel symbols selected from a predetermined signal constellation in response to said first and second encoder output streams, said receiver apparatus

CHARACTERIZED BY

This Sheet Has Been Replaced

* C L M S *

means (220) including a maximum likelihood decoder for recovering from said received signal a) said first data stream and b) said second encoder output stream,

a Reed-Solomon decoder (230), and

means (222) for applying said recovered second encoder output stream to said Reed-Solomon decoder to recover said second data stream.

6. A method CHARACTERIZED by the steps of

applying first and second streams of data to a trellis encoder (112) and a Reed-Solomon encoder (114), respectively, and

generating (on 150), in response to the outputs of said encoders, an output signal representing channel symbols selected from a predetermined signal constellation.

7. The invention of claim 6 wherein said signal constellation is comprised of a plurality of subsets of channel symbols, and wherein said generating step includes the steps of

identifying a sequence of said subsets in response to the output of said trellis encoder, and

selecting a channel symbol from each subset of said sequence in response to the output of said Reed-Solomon encoder.

8. The invention of claim 7 wherein said generating step includes the steps of

generating, as said output signal, a passband signal representing the selected channel symbols, and

applying said passband signal to a transmission channel.

9. The invention of claim 7 wherein the combination of said applying and generating steps provides a first particular level of error rate performance for said first stream of data and a second particular level of error-rate performance for said second stream of data which is at least as great as said first particular level of error rate performance.

10. The invention of claim 4 or 9 wherein said constellation is a 2N-dimensional constellation based on a rectangular lattice, N being one of the

This Sheet Has Been Replaced

integers 1, 2, 4, 8 ..., and wherein there are 4N of said subsets.

11. The invention of claim 5 or 10 wherein N is 2, and said Reed-Solomon encoder implements a single-error-correcting Reed-Solomon code.

12. The invention of claim 6 or 11 wherein said trellis encoder implements a four-dimensional, 64-state trellis encoder.

13. The invention of claim 5 or 10 wherein N is 1, and said Reed-Solomon encoder implements a double-error-correcting Reed-Solomon code.

14. The invention of claim 8 or 13 wherein said trellis encoder implements a two-dimensional, 64-state trellis encoder.

This Sheet Has Been Replaced