7 CHANNEL CODING: PART 2 381

7.1 Convolutional Encoding, 382
7.2 Convolutional Encoder Representation, 384
7.2.1 Connection Representation, 385
7.2.2 State Representation and the State Diagram, 389
7.2.3 The Tree Diagram, 391
7.2.4 The Trellis Diagram, 393
7.3 Formulation of the Convolutional Decoding Problem, 395
7.3.1 Maximum Likelihood Decoding, 395
6.4.9 Decoder Implementation, 340

7.3.2 Channel Models: Hard versus Soft Decisions, 396
7.3.3 The Viterbi Convolutional Decoding Algorithm, 401
7.3.4 An Example of Viterbi Convolutional Decoding, 401
7.3.5 Decoder Implementation, 405
7.3.6 Path Memory and Synchronization, 408
7.4 Properties of Convolutional Codes, 408
7.4.1 Distance Properties of Convolutional Codes, 408
7.4.2 Systematic and Nonsystematic Convolutional Codes, 413
7.4.3 Catastrophic Error Propagation in Convolutional Codes, 414
7.4.4 Performance Bounds for Convolutional Codes, 415
7.4.5 Coding Gain, 416
7.4.6 Best Known Convolutional Codes, 418
7.4.7 Convolutional Code Rate Trade-Off, 420
7.4.8 Soft-Decision Viterbi Decoding, 420
7.5 Other Convolutional Decoding Algorithms, 422
7.5.1 Sequential Decoding, 422
7.5.2 Comparisons and Limitations of Viterbi and Sequential Decoding, 425
7.5.3 Feedback Decoding, 427
7.6 Conclusion, 429

Channel Coding: Part 2

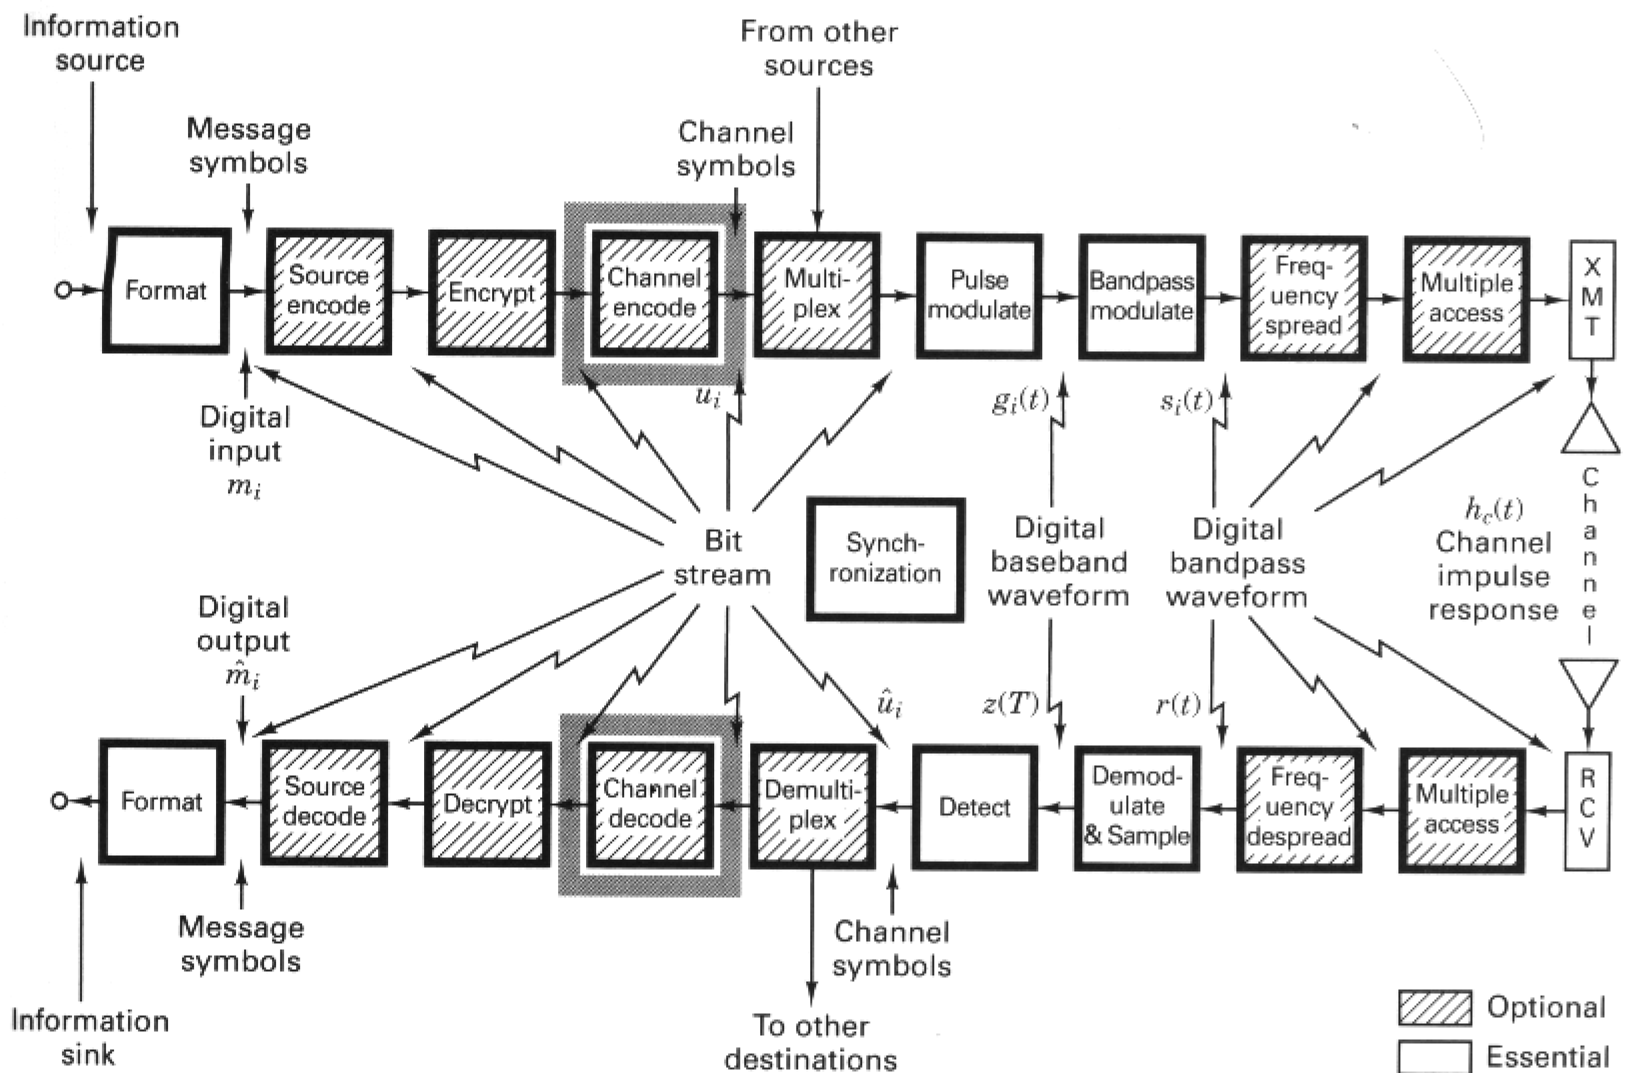

This chapter deals with convolutional coding. Chapter 6 presented the fundamentals of linear block codes, which are described by two integers, n and k, and a generator matrix or polynomial. The integer k is the number of data bits that form an input to a block encoder. The integer n is the total number of bits in the associated codeword out of the encoder. A characteristic of linear block codes is that each codeword n-tuple is uniquely determined by the input message k-tuple. The ratio k/n is called the *rate* of the code—a measure of the amount of added redundancy. A *convolutional code* is described by three integers, n, k, and K, where the ratio k/n has the same code rate significance (information per coded bit) that it has for block codes; however, n does *not* define a block or codeword length as it does for block codes. The integer K is a parameter known as the *constraint length*; it represents the number of k-tuple stages in the encoding shift register. An important characteristic of convolutional codes, different from block codes, is that the encoder has memory—the n-tuple emitted by the convolutional encoding procedure is not only a function of an input k-tuple, but is also a function of the previous $K - 1$ input k-tuples. In practice, n and k are small integers and K is varied to control the capability and complexity of the code.

7.1 CONVOLUTIONAL ENCODING

In Figure 1.2 we presented a typical block diagram of a digital communication system. A version of this functional diagram, focusing primarily on the convolutional encode/decode and modulate/demodulate portions of the communication link, is

shown in Figure 7.1. The input message source is denoted by the sequence $\mathbf{m} = m_1, m_2, \ldots, m_i, \ldots$, where each m_i represents a binary digit (bit), and i is a time index. To be precise, one should denote the elements of $\mathbf{m}$ with an index for class membership (e.g., for binary codes, 1 or 0) and an index for time. However, in this chapter, for simplicity, indexing is only used to indicate time (or location within a sequence). We shall assume that each m_i is equally likely to be a one or a zero, and independent from digit to digit. Being independent, the bit sequence lacks any redundancy; that is, knowledge about bit m_i gives no information about m_j $(i \neq j)$. The encoder transforms each sequence $\mathbf{m}$ into a unique codeword sequence $\mathbf{U} = G(\mathbf{m})$. Even though the sequence $\mathbf{m}$ uniquely defines the sequence $\mathbf{U}$, a key feature of convolutional codes is that a given k-tuple within $\mathbf{m}$ does *not* uniquely define its associated n-tuple within $\mathbf{U}$ since the encoding of each k-tuple is *not only* a function of that k-tuple but is also a function of the $K - 1$ input k-tuples that precede it. The sequence $\mathbf{U}$ can be partitioned into a sequence of branch words: $\mathbf{U} = U_1, U_2, \ldots, U_i, \ldots$. Each branch word U_i is made up of binary *code symbols*, often called *channel symbols, channel bits,* or *code bits;* unlike the input message bits the code symbols are not independent.

In a typical communication application, the codeword sequence $\mathbf{U}$ modulates a waveform $s(t)$. During transmission, the waveform $s(t)$ is corrupted by noise, resulting in a received waveform $\hat{s}(t)$ and a demodulated sequence $\mathbf{Z} = Z_1, Z_2, \ldots, Z_i, \ldots$, as indicated in Figure 7.1. The task of the decoder is to produce an estimate $\hat{\mathbf{m}} = \hat{m}_1, \hat{m}_2, \ldots, \hat{m}_i, \ldots$ of the original message sequence, using the received sequence $\mathbf{Z}$ together with a priori knowledge of the encoding procedure.

A general convolutional encoder, shown in Figure 7.2, is mechanized with a kK-stage shift register and n modulo-2 adders, where K is the constraint length.

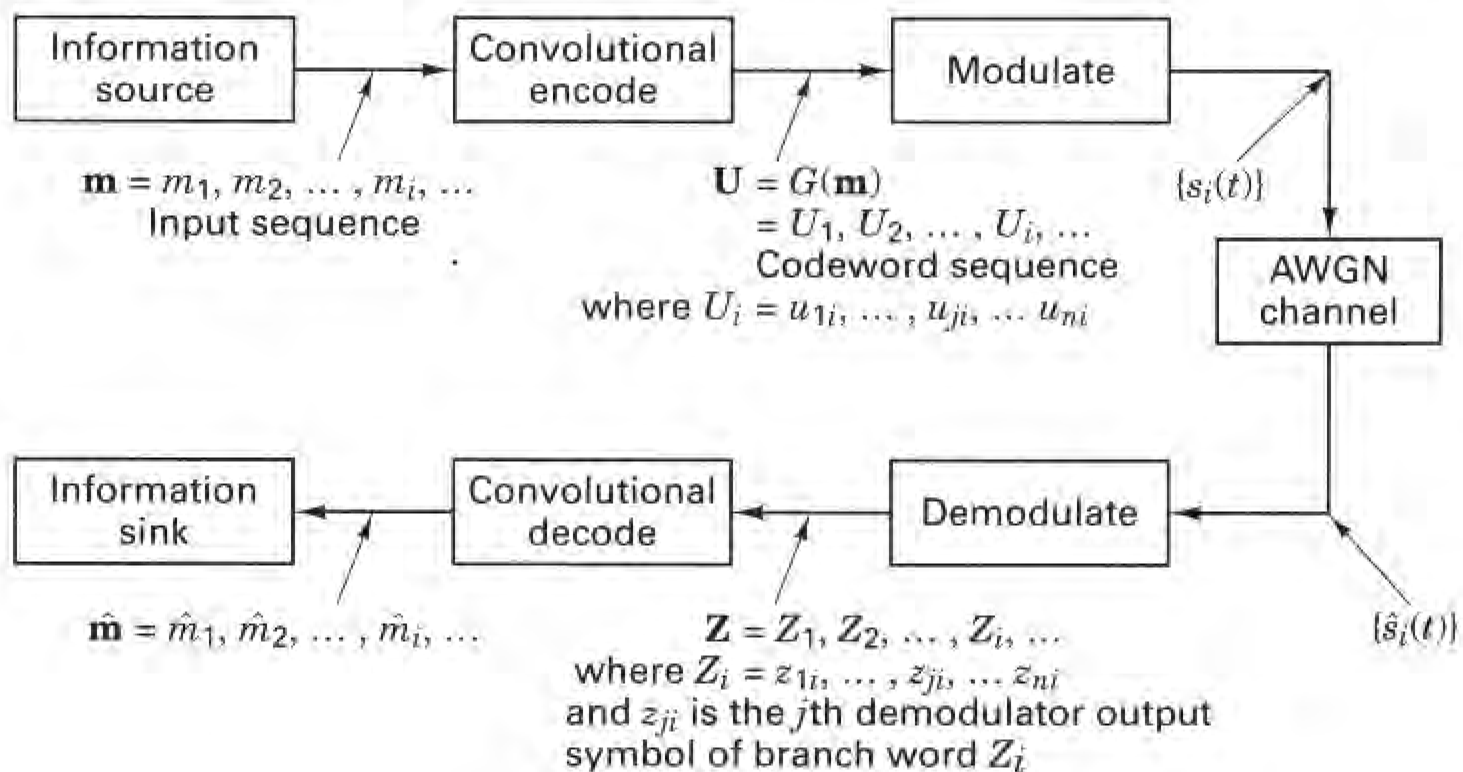

Figure 7.1 Encode/decode and modulate/demodulate portions of a communication link.

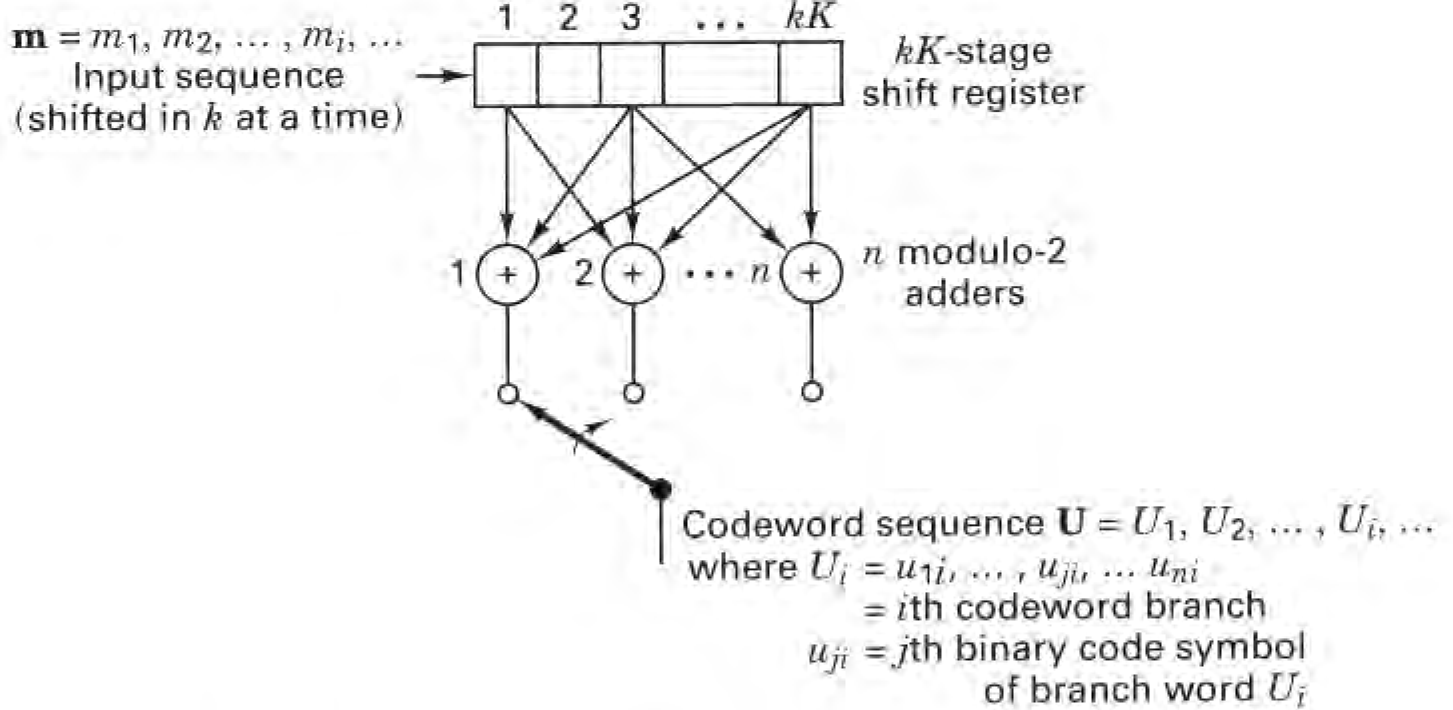

Figure 7.2 Convolutional encoder with constraint length K and rate k/n.

The constraint length represents the number of k-bit shifts over which a single information bit can influence the encoder output. At each unit of time, k bits are shifted into the first k stages of the register; all bits in the register are shifted k stages to the right, and the outputs of the n adders are sequentially sampled to yield the binary code symbols or code bits. These code symbols are then used by the modulator to specify the waveforms to be transmitted over the channel. Since there are n code bits for each input group of k message bits, the code rate is k/n message bit per code bit, where $k < n$.

We shall consider only the most commonly used binary convolutional encoders for which $k = 1$—that is, those encoders in which the message bits are shifted into the encoder one bit at a time, although generalization to higher order alphabets is straightforward [1, 2]. For the $k = 1$ encoder, at the ith unit of time, message bit m_i is shifted into the first shift register stage; all previous bits in the register are shifted one stage to the right, and as in the more general case, the outputs of the n adders are sequentially sampled and transmitted. Since there are n code bits for each message bit, the code rate is $1/n$. The n code symbols occurring at time t_i comprise the ith branch word, $U_i = u_{1i}, u_{2i}, \dots, u_{ni}$, where u_{ji} ($j = 1, 2, \dots, n$) is the jth code symbol belonging to the ith branch word. Note that for the rate $1/n$ encoder, the kK-stage shift register can be referred to simply as a K-stage register, and the constraint length K, which was expressed in units of k-tuple stages, can be referred to as constraint length in units of bits.

7.2 CONVOLUTIONAL ENCODER REPRESENTATION

To describe a convolutional code, one needs to characterize the encoding function $G(\mathbf{m})$, so that given an input sequence $\mathbf{m}$, one can readily compute the output sequence $\mathbf{U}$. Several methods are used for representing a convolutional encoder, the

most popular being the *connection pictorial, connection vectors or polynomials,* the *state diagram,* the *tree diagram,* and the *trellis diagram.* They are each described below.

7.2.1 Connection Representation

We shall use the convolutional encoder, shown in Figure 7.3, as a model for discussing convolutional encoders. The figure illustrates a (2, 1) convolutional encoder with constraint length $K = 3$. There are $n = 2$ modulo-2 adders; thus the code rate k/n is $\frac{1}{2}$. At each input bit time, a bit is shifted into the leftmost stage and the bits in the register are shifted one position to the right. Next, the output switch samples the output of each modulo-2 adder (i.e., first the upper adder, then the lower adder), thus forming the code symbol pair making up the branch word associated with the bit just inputted. The sampling is repeated for each inputted bit. The choice of connections between the adders and the stages of the register gives rise to the characteristics of the code. Any change in the choice of connections results in a different code. The connections are, of course, *not* chosen or changed arbitrarily. The problem of choosing connections to yield good distance properties is complicated and has not been solved in general; however, good codes have been found by computer search for all constraint lengths less than about 20 [3–5].

Unlike a block code that has a fixed word length n, a convolutional code has no particular block size. However, convolutional codes are often forced into a block structure by *periodic truncation.* This requires a number of zero bits to be appended to the end of the input data sequence, for the purpose of clearing or *flushing* the encoding shift register of the data bits. Since the added zeros carry no information, the *effective code rate* falls below k/n. To keep the code rate close to k/n, the truncation period is generally made as long as practical.

One way to represent the encoder is to specify a set of n *connection vectors,* one for each of the n modulo-2 adders. Each vector has dimension K and describes the connection of the encoding shift register to that modulo-2 adder. A one in the ith position of the vector indicates that the corresponding stage in the shift register

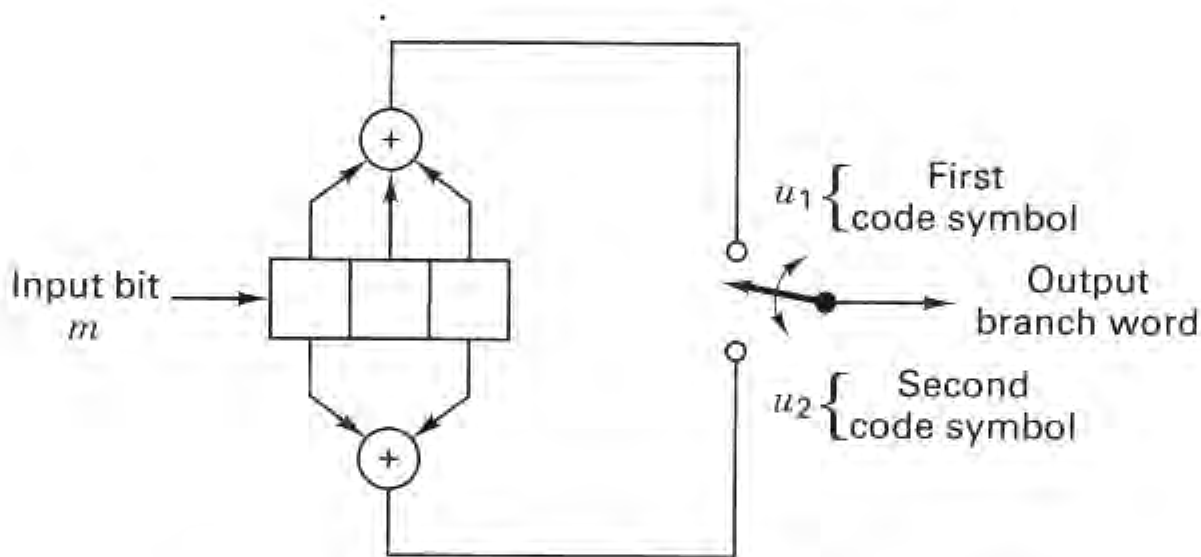

Figure 7.3 Convolutional encoder (rate $\frac{1}{2}$, $K = 3$).

is connected to the modulo-2 adder, and a zero in a given position indicates that no connection exists between the stage and the modulo-2 adder. For the encoder example in Figure 6.3, we can write the connection vector $\mathbf{g}_1$ for the upper connections and $\mathbf{g}_2$ for the lower connections as follows:

$$\mathbf{g}_1 = 1\ 1\ 1$$
$$\mathbf{g}_2 = 1\ 0\ 1$$

Now consider that a message vector $\mathbf{m} = 1\ 0\ 1$ is convolutionally encoded with the encoder shown in Figure 7.3. The three message bits are inputted, one at a time, at times t_1, t_2, and t_3, as shown in Figure 7.4. Subsequently, $(K - 1) = 2$ zeros are inputted at times t_4 and t_5 to flush the register and thus ensure that the tail end of the message is shifted the full length of the register. The output sequence is seen to be 1 1 1 0 0 0 1 0 1 1, where the leftmost symbol represents the earliest transmission. The entire output sequence, including the code symbols as a result of flushing, are needed to decode the message. To flush the message from the encoder requires one less zero than the number of stages in the register, or $K - 1$ flush bits. Another zero input is shown at time t_6, for the reader to verify that the flushing is completed at time t_5. Thus, a new message can be entered at time t_6.

7.2.1.1 Impulse Response of the Encoder

We can approach the encoder in terms of its *impulse response*—that is, the response of the encoder to a single "one" bit that moves through it. Consider the contents of the register in Figure 7.3 as a one moves through it:

Register contents	Branch word	
	u_1	u_2
100	1	1
010	1	0
001	1	1

Input sequence: 1 0 0
Output sequence: 11 10 11

The output sequence for the input "one" is called the impulse response of the encoder. Then, for the input sequence $\mathbf{m} = 1\ 0\ 1$, the output may be found by the *superposition* or the *linear addition* of the time-shifted input "impulses" as follows:

Input **m**	Output				
1	11	10	11		
0		00	00	00	
1			11	10	11
Modulo-2 sum:	11	10	00	10	11

Observe that this is the same output as that obtained in Figure 7.4, demonstrating that *convolutional codes are linear*—just like the linear block codes of Chapter 6. It

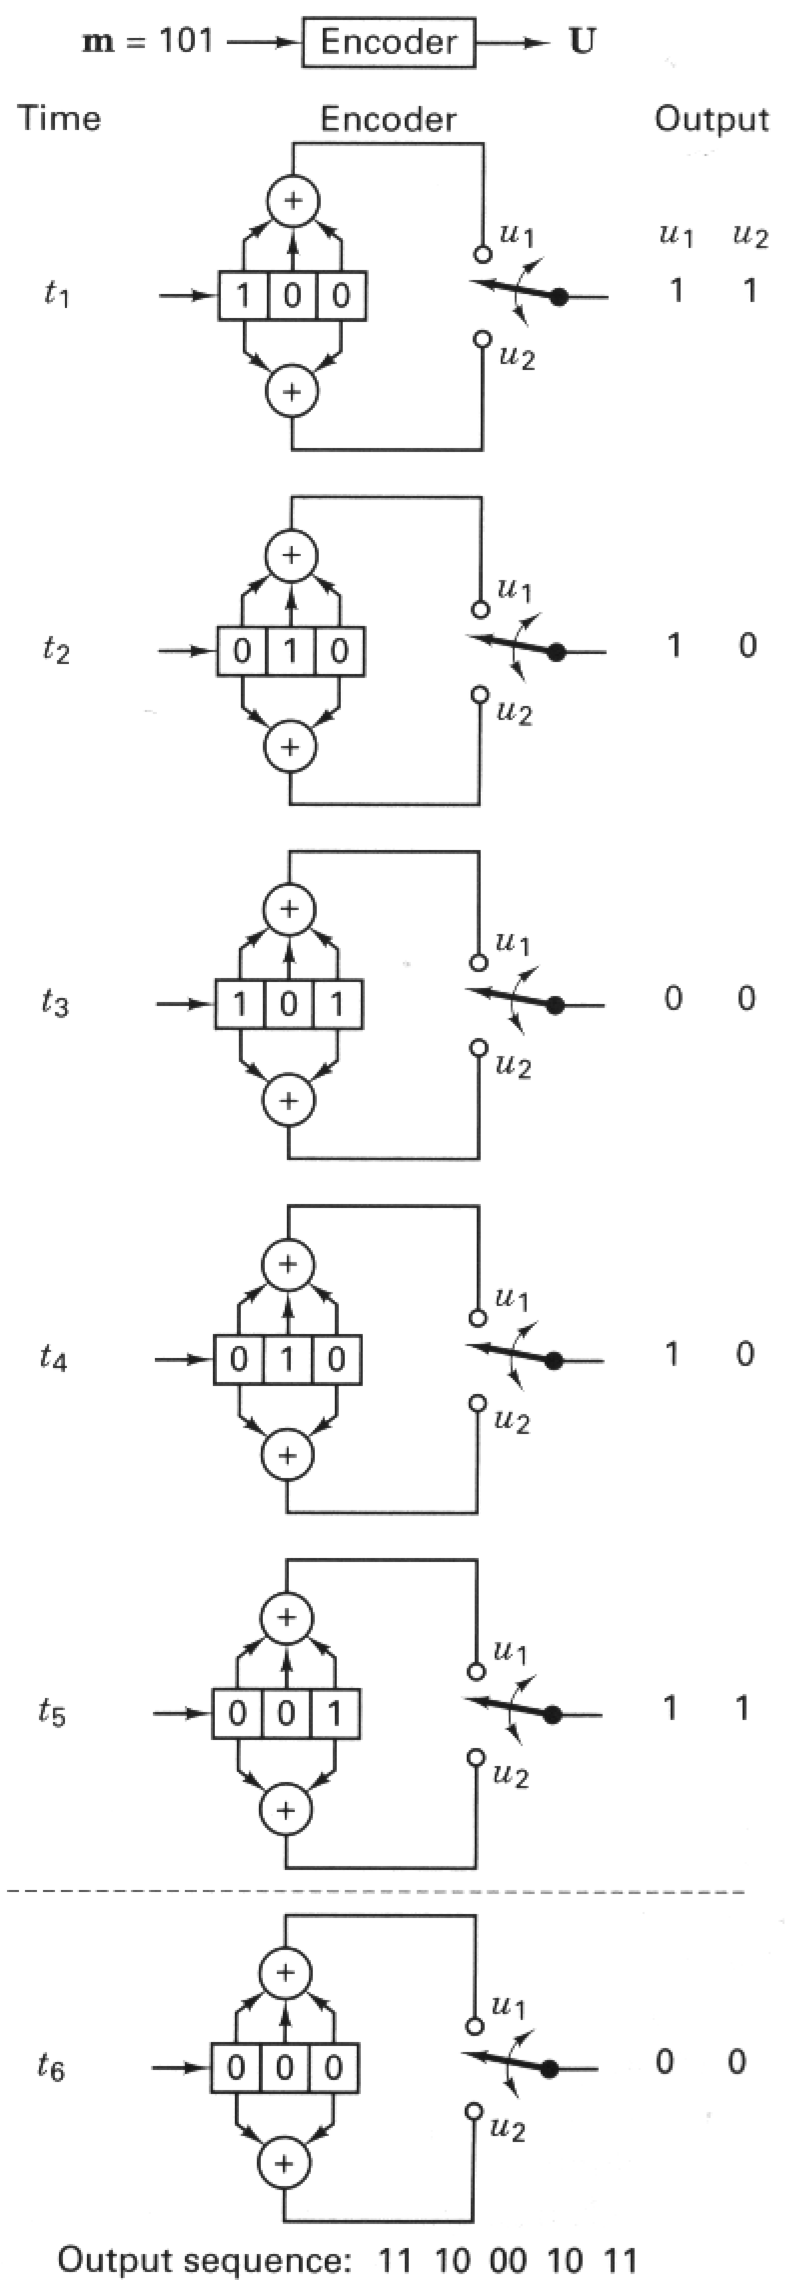

Figure 7.4 Convolutionally encoding a message sequence with a rate $\frac{1}{2}$, $K = 3$ encoder.

is from this property of generating the output by the linear addition of time-shifted impulses, or the convolution of the input sequence with the impulse response of the encoder, that we derive the name *convolutional encoder.* Often, this encoder characterization is presented in terms of an infinite-order generator matrix [6].

Notice that the *effective code rate* for the foregoing example with 3-bit input sequence and 10-bit output sequence is $k/n = \frac{3}{10}$—quite a bit less than the rate $\frac{1}{2}$ that might have been expected from the knowledge that each input data bit yields a pair of output channel bits. The reason for the disparity is that the final data bit into the encoder needs to be shifted through the encoder. All of the output channel bits are needed in the decoding process. If the message had been longer, say 300 bits, the output codeword sequence would contain 604 bits, resulting in a code rate of 300/604—much closer to $\frac{1}{2}$.

7.2.1.2 Polynomial Representation

Sometimes, the encoder connections are characterized by *generator polynomials,* similar to those used in Chapter 6 for describing the feedback shift register implementation of cyclic codes. We can represent a convolutional encoder with a set of n generator polynomials, one for each of the n modulo-2 adders. Each polynomial is of degree $K - 1$ or less and describes the connection of the encoding shift register to that modulo-2 adder, much the same way that a connection vector does. The coefficient of each term in the $(K - 1)$-degree polynomial is either 1 or 0, depending on whether a connection exists or does not exist between the shift register and the modulo-2 adder in question. For the encoder example in Figure 7.3, we can write the generator polynomial $\mathbf{g}_1(X)$ for the upper connections and $\mathbf{g}_2(X)$ for the lower connections as follows:

$$\mathbf{g}_1(X) = 1 + X + X^2$$
$$\mathbf{g}_2(X) = 1 + X^2$$

where the lowest order term in the polynomial corresponds to the input stage of the register. The output sequence is found as follows:

$$\mathbf{U}(X) = \mathbf{m}(X)\mathbf{g}_1(X) \text{ interlaced with } \mathbf{m}(X)\mathbf{g}_2(X)$$

First, express the message vector $\mathbf{m} = 1\ 0\ 1$ as a polynomial—that is, $\mathbf{m}(X) = 1 + X^2$. We shall again assume the use of zeros following the message bits, to flush the register. Then the output polynomial $\mathbf{U}(X)$, or the output sequence $\mathbf{U}$, of the Figure 7.3 encoder can be found for the input message $\mathbf{m}$ as follows:

$$\mathbf{m}(X)\mathbf{g}_1(X) = (1 + X^2)(1 + X + X^2) = 1 + X + X^3 + X^4$$
$$\mathbf{m}(X)\mathbf{g}_2(X) = (1 + X^2)(1 + X^2) = 1 + X^4$$

$$\mathbf{m}(X)\mathbf{g}_1(X) = 1 + X + 0X^2 + X^3 + X^4$$
$$\mathbf{m}(X)\mathbf{g}_2(X) = 1 + 0X + 0X^2 + 0X^3 + X^4$$

$$\mathbf{U}(X) = (1, 1) + (1, 0)X + (0, 0)X^2 + (1, 0)X^3 + (1, 1)X^4$$
$$\mathbf{U} = 1\,1 \quad 1\,0 \quad 0\,0 \quad 1\,0 \quad 1\,1$$

In this example we started with another point of view—namely, that the convolutional encoder can be treated as a set of *cyclic code shift registers.* We represented the encoder with *polynomial generators* as used for describing cyclic codes. However, we arrived at the same output sequence as in Figure 7.4 and at the same output sequence as the impulse response treatment of the preceding section. (For a good presentation of convolutional code structure in the context of linear sequential circuits, see Reference [7].)

7.2.2 State Representation and the State Diagram

A convolutional encoder belongs to a class of devices known as *finite-state machines,* which is the general name given to machines that have a memory of past signals. The adjective *finite* refers to the fact that there are only a finite number of unique states that the machine can encounter. What is meant by the *state* of a finite-state machine? In the most general sense, the state consists of the smallest amount of information that, together with a current input to the machine, can predict the output of the machine. The state provides some knowledge of the past signaling events and the restricted set of possible outputs in the future. A future state is restricted by the past state. For a rate $1/n$ convolutional encoder, the state is represented by the contents of the rightmost $K-1$ stages (see Figure 7.3). Knowledge of the state together with knowledge of the next input is necessary and sufficient to determine the next output. Let the state of the encoder at time t_i be defined as $X_i = m_{i-1}, m_{i-2}, \ldots, m_{i-K+1}$. The ith codeword branch U_i is completely determined by state X_i and the present input bit m_i; thus the state X_i represents the past history of the encoder in determining the encoder output. The encoder state is said to be *Markov,* in the sense that the probability $P(X_{i+1}|X_i, X_{i-1}, \ldots, X_0)$ of being in state X_{i+1}, given all previous states, depends only on the most recent state X_i; that is, the probability is equal to $P(X_{i+1}|X_i)$.

One way to represent simple encoders is with a *state diagram;* such a representation for the encoder in Figure 7.3 is shown in Figure 7.5. The states, shown in the boxes of the diagram, represent the possible contents of the rightmost $K-1$ stages of the register, and the paths between the states represent the output branch words resulting from such state transitions. The states of the register are designated $a = 00$, $b = 10$, $c = 01$, and $d = 11$; the diagram shown in Figure 7.5 illustrates all the state transitions that are possible for the encoder in Figure 7.3. There are *only two transitions* emanating from each state, corresponding to the two possible input bits. Next to each path between states is written the output branch word associated with the state transition. In drawing the path, we use the convention that a solid line denotes a path associated with an input bit, zero, and a dashed line denotes a path associated with an input bit, one. Notice that it is *not possible* in a single transition to move from a given state to *any arbitrary state.* As a consequence of shifting-in one bit at a time, there are only two possible state transitions that the register can make at each bit time. For example, if the present encoder state is 00, the *only possibilities* for the state at the next shift are 00 or 10.

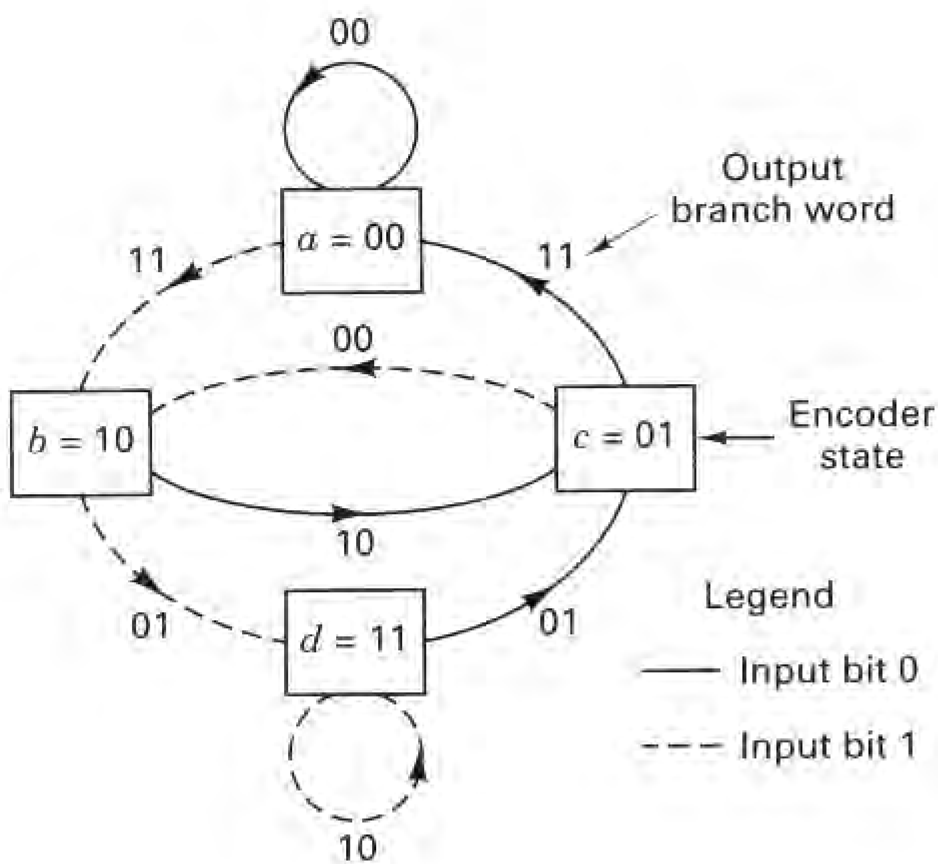

Figure 7.5 Encoder state diagram (rate $\frac{1}{2}$, $K = 3$).

Example 7.1 Convolutional Encoding

For the encoder shown in Figure 7.3, show the state changes and the resulting output codeword sequence **U** for the message sequence **m** = 1 1 0 1 1, followed by $K - 1 = 2$ zeros to flush the register. Assume that the initial contents of the register are all zeros.

Solution

Input bit m_i	Register contents	State at time t_i	State at time t_{i+1}	Branch word at time t_i u_1	u_2
—	0 0 0	0 0	0 0	—	
1	1 0 0	0 0	1 0	1	1
1	1 1 0	1 0	1 1	0	1
0	0 1 1	1 1	0 1	0	1
1	1 0 1	0 1	1 0	0	0
1	1 1 0	1 0	1 1	0	1
0	0 1 1	1 1	0 1	0	1
0	0 0 1	0 1	0 0	1	1

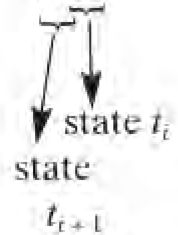

Output sequence: **U** = 1 1 0 1 0 1 0 0 0 1 0 1 1 1

Example 7.2 Convolutional Encoding

In Example 7.1 the initial contents of the register are all zeros. This is equivalent to the condition that the given input sequence is preceded by two zero bits (the encoding is a function of the present bit and the $K - 1$ prior bits). Repeat Example 7.1 with the assumption that the given input sequence is preceded by two one bits, and verify that now the codeword sequence $\mathbf{U}$ for input sequence $\mathbf{m} = 1\ 1\ 0\ 1\ 1$ is different than the codeword found in Example 7.1.

Solution

The entry "×" signifies "don't know."

Input bit m_i	Register contents	State at time t_i	State at time t_{i+1}	Branch word at time t_i u_1	u_2
—	1 1 ×	1 ×	1 1	—	
1	1 1 1	1 1	1 1	1	0
1	1 1 1	1 1	1 1	1	0
0	0 1 1	1 1	0 1	0	1
1	1 0 1	0 1	1 0	0	0
1	1 1 0	1 0	1 1	0	1
0	0 1 1	1 1	0 1	0	1
0	0 0 1	0 1	0 0	1	1

Output sequence: $\mathbf{U} = 1\,0 \quad 1\,0 \quad 0\,1 \quad 0\,0 \quad 0\,1 \quad 0\,1 \quad 1\,1$

By comparing this result with that of Example 7.1, we can see that each branch word of the output sequence $\mathbf{U}$ is *not only* a function of the input bit, but is also a function of the $K - 1$ prior bits.

7.2.3 The Tree Diagram

Although the state diagram completely characterizes the encoder, one cannot easily use it for tracking the encoder transitions as a function of time since the diagram cannot represent time history. The tree diagram adds the *dimension of time* to the state diagram. The tree diagram for the convolutional encoder shown in Figure 7.3 is illustrated in Figure 7.6. At each successive input bit time the encoding procedure can be described by traversing the diagram from left to right, each tree branch describing an output branch word. The branching rule for finding a codeword sequence is as follows: If the input bit is a zero, its associated branch word is found by moving to the next rightmost branch in the upward direction. If the input bit is a one, its branch word is found by moving to the next rightmost branch in the downward direction. Assuming that the initial contents of the encoder is all zeros, the

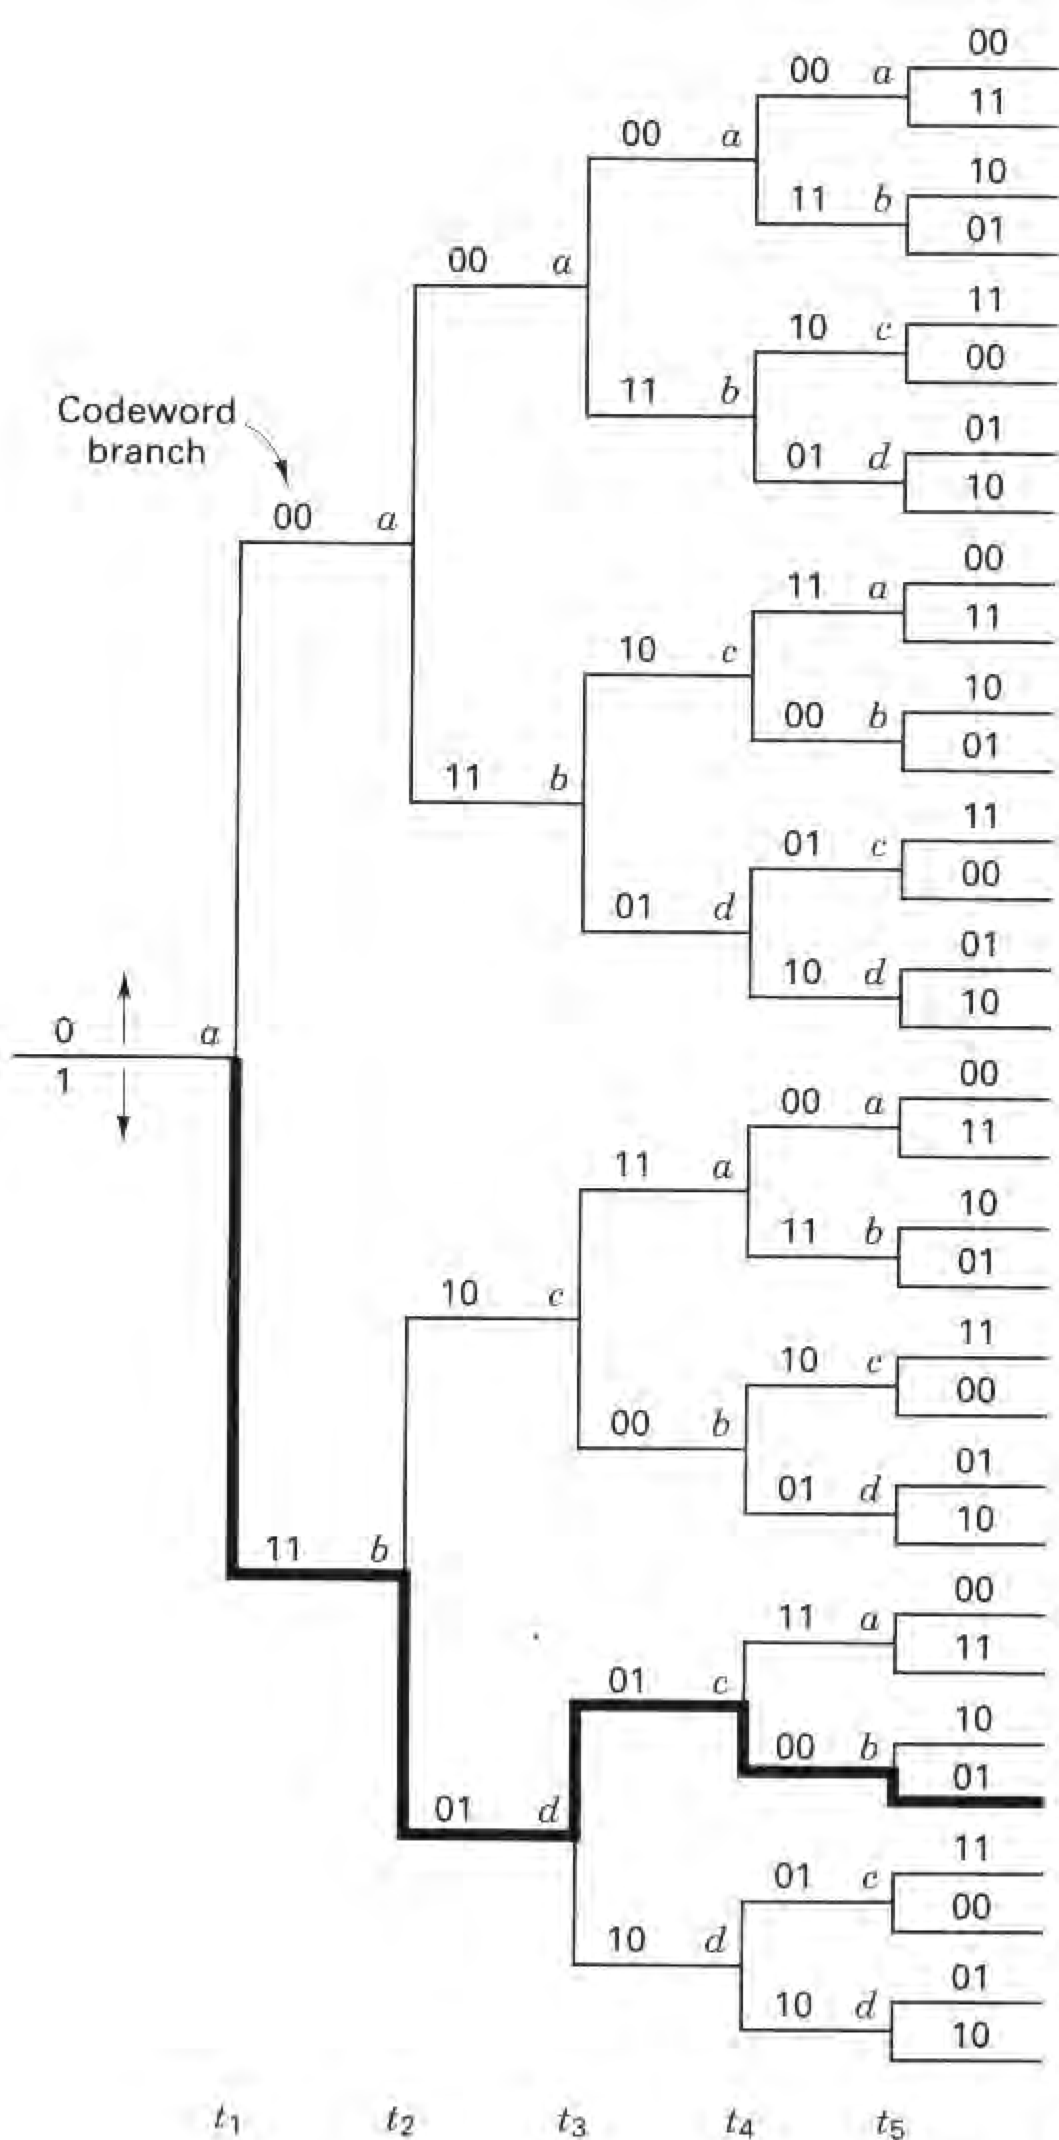

Figure 7.6 Tree representation of encoder (rate $\frac{1}{2}$, $K = 3$).

diagram shows that if the first input bit is a zero, the output branch word is 00 and, if the first input bit is a one, the output branch word is 11. Similarly, if the first input bit is a one and the second input bit is a zero, the second output branch word is 10. Or, if the first input bit is a one and the second input bit is a one, the second output branch word is 01. Following this procedure we see that the input sequence 1 1 0 1 1 traces the heavy line drawn on the tree diagram in Figure 7.6. This path corresponds to the output codeword sequence 1 1 0 1 0 1 0 0 0 1.

The added dimension of time in the tree diagram (compared to the state diagram) allows one to dynamically describe the encoder as a function of a particular input sequence. However, can you see one problem in trying to use a tree diagram for describing a sequence of any length? The number of branches increases as a function of 2^L, where L is the number of branch words in the sequence. You would quickly run out of paper, and patience.

7.2.4 The Trellis Diagram

Observation of the Figure 7.6 tree diagram shows that for this example, the structure repeats itself at time t_4, after the third branching (in general, the tree structure *repeats after K branchings,* where K is the constraint length). We label each node in the tree of Figure 7.6 to correspond to the four possible states in the shift register, as follows: $a = 00$, $b = 10$, $c = 01$, and $d = 11$. The first branching of the tree structure, at time t_1, produces a pair of nodes labeled a and b. At each successive branching the number of nodes double. The second branching, at time t_2, results in four nodes labeled *a, b, c,* and *d.* After the *third* branching, there are a total of eight nodes: two are labeled *a,* two are labeled *b,* two are labeled *c,* and two are labeled *d.* We can see that all branches emanating from two nodes of the same state generate identical branch word sequences. From this point on, the upper and the lower halves of the tree are identical. The reason for this should be obvious from examination of the encoder in Figure 7.3. As the fourth input bit enters the encoder on the left, the first input bit is ejected on the right and no longer influences the output branch words. Consequently, the input sequences 1 0 0 x y . . . and 0 0 0 x y . . . , where the leftmost bit is the earliest bit, generate the same branch words after the ($K = 3$)rd branching. This means that any two nodes having the same state label at the same time t_i can be merged, since all succeeding paths will be indistinguishable. If we do this to the tree structure of Figure 7.6, we obtain another diagram, called the trellis diagram. The *trellis diagram,* by exploiting the repetitive structure, provides a more manageable encoder description than does the tree diagram. The trellis diagram for the convolutional encoder of Figure 7.3 is shown in Figure 7.7.

In drawing the trellis diagram, we use the same convention that we introduced with the state diagram—a solid line denotes the output generated by an input bit zero, and a dashed line denotes the output generated by an input bit one. The nodes of the trellis characterize the encoder states; the first row nodes correspond to the state $a = 00$, the second and subsequent rows correspond to the states $b = 10$, $c = 01$, and $d = 11$. At each unit of time, the trellis requires 2^{K-1} nodes to represent the 2^{K-1} possible encoder states. The trellis in our example assumes a

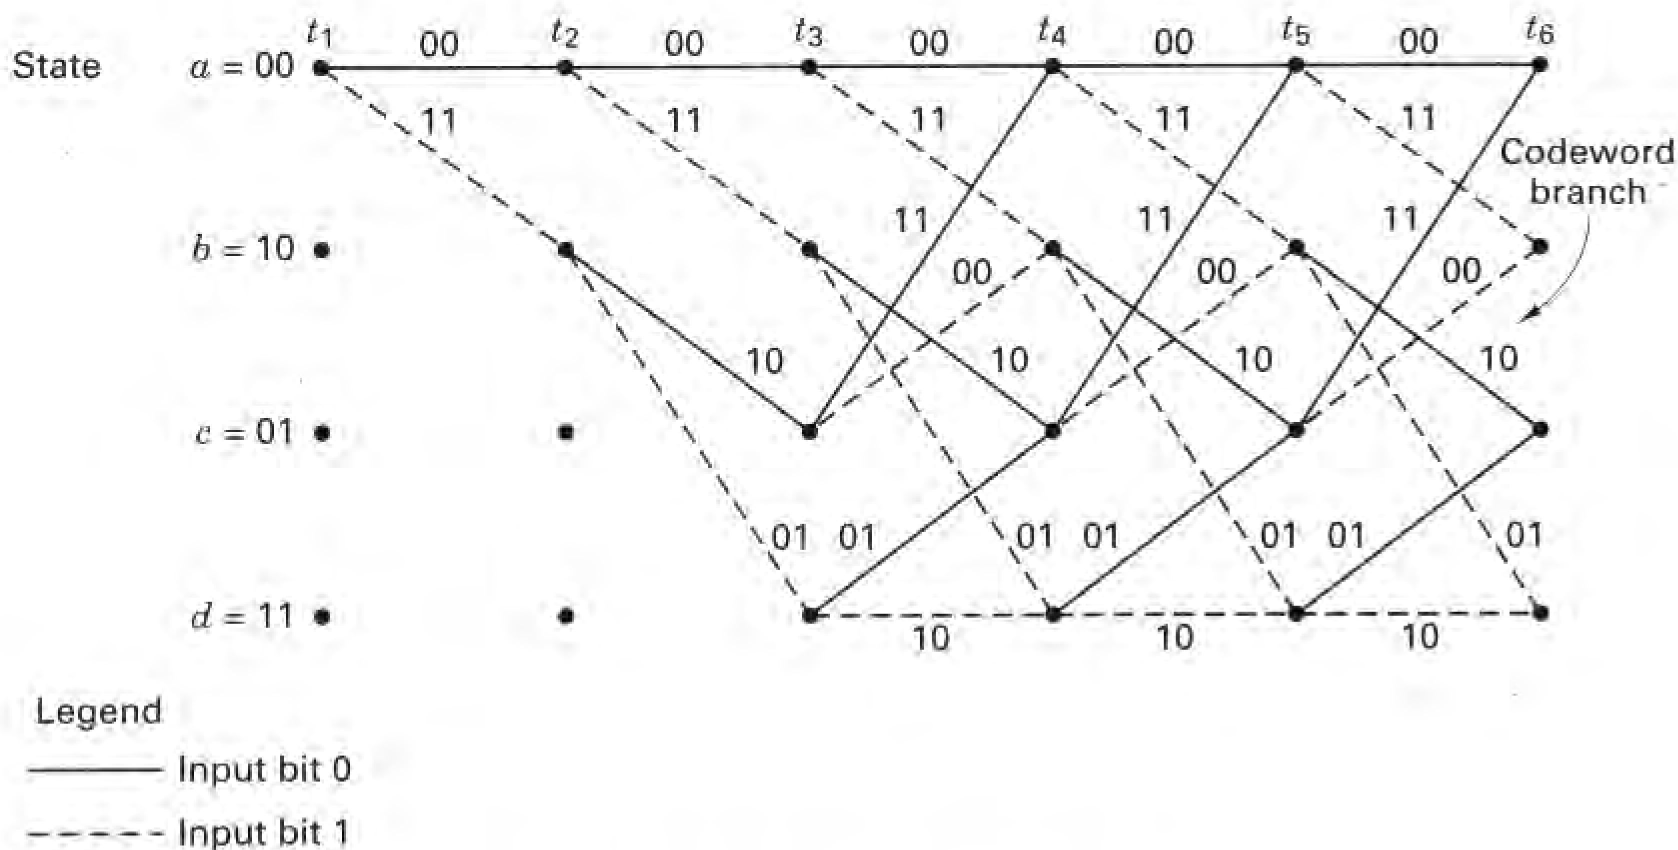

Figure 7.7 Encoder trellis diagram (rate $\frac{1}{2}$, $K = 3$).

fixed periodic structure after trellis depth 3 is reached (at time t_4). In the general case, the fixed structure prevails after depth K is reached. At this point and thereafter, each of the states can be entered from either of two preceding states. Also, each of the states can transition to one of two states. Of the two outgoing branches, one corresponds to an input bit zero and the other corresponds to an input bit one. On Figure 7.7 the output branch words corresponding to the state transitions appear as labels on the trellis branches.

One time-interval section of a fully-formed encoding trellis structure completely defines the code. The only reason for showing several sections is for viewing a code-symbol sequence as a function of time. The state of the convolutional encoder is represented by the contents of the rightmost $K - 1$ stages in the encoder register. Some authors describe the state as the contents of the leftmost $K - 1$ stages. Which description is correct? They are both correct in the following sense. Every transition has a starting state and a terminating state. The rightmost $K - 1$ stages describe the starting state for the current input, which is in the leftmost stage (assuming a rate $1/n$ encoder). The leftmost $K - 1$ stages represent the terminating state for that transition. A code-symbol sequence is characterized by N branches (representing N data bits) occupying N intervals of time and associated with a particular state at each of $N + 1$ times (from start to finish). Thus, we launch bits at times $t_1, t_2, \ldots, t_N$, and are interested in state metrics at times $t_1, t_2, \ldots, t_{N+1}$. The convention used here is that the current bit is located in the leftmost stage (not on a wire leading to that stage), and the rightmost $K - 1$ stages start in the all-zeros state. We refer to this time as the *start time* and label it t_1. We refer to the concluding time of the last transition as the *terminating time* and label it t_{N+1}.

7.3 FORMULATION OF THE CONVOLUTIONAL DECODING PROBLEM

7.3.1 Maximum Likelihood Decoding

If all input message sequences are equally likely, a decoder that achieves the minimum probability of error is one that compares the conditional probabilities, also called the *likelihood functions* $P(\mathbf{Z}|\mathbf{U}^{(m)})$, where $\mathbf{Z}$ is the received sequence and $\mathbf{U}^{(m)}$ is one of the possible transmitted sequences, and chooses the maximum. The decoder chooses $\mathbf{U}^{(m')}$ if

$$P(\mathbf{Z}|\mathbf{U}^{(m')}) = \max_{\text{over all } \mathbf{U}^{(m)}} P(\mathbf{Z}|\mathbf{U}^{(m)}) \tag{7.1}$$

The *maximum likelihood* concept, as stated in Equation (7.1), is a fundamental development of decision theory (see Appendix B); it is the formalization of a "common-sense" way to make decisions when there is statistical knowledge of the possibilities. In the binary demodulation treatment in Chapters 3 and 4 there were *only two* equally likely possible signals, $s_1(t)$ or $s_2(t)$, that might have been transmitted. Therefore, to make the binary maximum likelihood decision, given a received signal, meant only to decide that $s_1(t)$ was transmitted if

$$p(z|s_1) > p(z|s_2)$$

otherwise, to decide that $s_2(t)$ was transmitted. The parameter z represents $z(T)$, the receiver predetection value at the end of each symbol duration time $t = T$. However, when applying maximum likelihood to the convolutional decoding problem, we observe that the convolutional code has memory (the received sequence represents the superposition of current bits and prior bits). Thus, applying maximum likelihood to the decoding of convolutionally encoded bits is performed in the context of choosing the *most likely sequence,* as shown in Equation (7.1). There are typically a *multitude* of possible codeword sequences that might have been transmitted. To be specific, for a binary code, a sequence of L branch words is a member of a set of 2^L possible sequences. Therefore, in the maximum likelihood context, we can say that the decoder chooses a particular $\mathbf{U}^{(m')}$ as the transmitted sequence if the likelihood $P(\mathbf{Z}|\mathbf{U}^{(m')})$ is greater than the likelihoods of all the other possible transmitted sequences. Such an optimal decoder, which minimizes the error probability (for the case where all transmitted sequences are equally likely), is known as a *maximum likelihood decoder.* The likelihood functions are given or computed from the specifications of the channel.

We will assume that the noise is additive white Gaussian with zero mean and the channel is *memoryless,* which means that the noise affects each code symbol *independently* of all the other symbols. For a convolutional code of rate $1/n$, we can therefore express the likelihood as

$$P(\mathbf{Z}|\mathbf{U}^{(m)}) = \prod_{i=1}^{\infty} P(Z_i|U_i^{(m)}) = \prod_{i=1}^{\infty} \prod_{j=1}^{n} P(z_{ji}|u_{ji}^{(m)}) \tag{7.2}$$

where Z_i is the ith branch of the received sequence $\mathbf{Z}$, $U_i^{(m)}$ is the ith branch of a particular codeword sequence $\mathbf{U}^{(m)}$, z_{ji} is the jth code symbol of Z_i, and $u_{ji}^{(m)}$ is the jth code symbol of $U_i^{(m)}$, and each branch comprises n code symbols. The decoder problem consists of choosing a path through the trellis of Figure 7.7 (each possible path defines a codeword sequence) such that

$$\prod_{i=1}^{\infty} \prod_{j=1}^{n} P(z_{ji}|u_{ji}^{(m)}) \text{ is maximized} \tag{7.3}$$

Generally, it is computationally more convenient to use the logarithm of the likelihood function since this permits the summation, instead of the multiplication, of terms. We are able to use this transformation because the logarithm is a monotonically increasing function and thus will not alter the final result in our codeword selection. We can define the log-likelihood function as

$$\gamma_{\mathbf{U}}(m) = \log P(\mathbf{Z}|\mathbf{U}^{(m)}) = \sum_{i=1}^{\infty} \log P(Z_i|U_i^{(m)}) = \sum_{i=1}^{\infty} \sum_{j=1}^{n} \log P(z_{ji}|u_{ji}^{(m)}) \tag{7.4}$$

The decoder problem now consists of choosing a path through the tree of Figure 7.6 or the trellis of Figure 7.7 such that $\gamma_{\mathbf{U}}(m)$ is maximized. For the decoding of convolutional codes, either the tree or the trellis structure can be used. In the tree representation of the code, the fact that the paths remerge is ignored. Since for a binary code, the number of possible sequences made up of L branch words is 2^L, maximum likelihood decoding of such a received sequence, using a tree diagram, requires the "brute force" or exhaustive comparison of 2^L accumulated log-likelihood metrics, representing all the possible different codeword sequences that could have been transmitted. Hence it is not practical to consider maximum likelihood decoding with a tree structure. It is shown in a later section that with the use of the trellis representation of the code, it is possible to configure a decoder which can discard the paths that could not possibly be candidates for the maximum likelihood sequence. The decoded path is chosen from some reduced set of *surviving paths.* Such a decoder is still optimum in the sense that the decoded path is the same as the decoded path obtained from a "brute force" maximum likelihood decoder, but the early rejection of unlikely paths reduces the decoding complexity.

For an excellent tutorial on the structure of convolutional codes, maximum likelihood decoding, and code performance, see Reference [8]. There are several algorithms that yield *approximate* solutions to the maximum likelihood decoding problem, including sequential [9, 10] and threshold [11]. Each of these algorithms is suited to certain special applications, but are all suboptimal. In contrast, the *Viterbi decoding algorithm* performs maximum likelihood decoding and is therefore optimal. This does not imply that the Viterbi algorithm is best for every application; there are severe constraints imposed by hardware complexity. The Viterbi algorithm is considered in Sections 7.3.3 and 7.3.4.

7.3.2 Channel Models: Hard versus Soft Decisions

Before specifying an algorithm that will determine the maximum likelihood decision, let us describe the channel. The codeword sequence $\mathbf{U}^{(m)}$, made up of branch

words, with each branch word comprised of n code symbols, can be considered to be an endless stream, as opposed to a block code, in which the source data and their codewords are partitioned into precise block sizes. The codeword sequence shown in Figure 7.1 emanates from the convolutional encoder and enters the modulator, where the code symbols are transformed into signal waveforms. The modulation may be baseband (e.g., pulse waveforms) or bandpass (e.g., PSK or FSK). In general, ℓ symbols at a time, where ℓ is an integer, are mapped into signal waveforms $s_i(t)$, where $i = 1, 2, \ldots, M = 2^\ell$. When $\ell = 1$, the modulator maps each code symbol into a binary waveform. The channel over which the waveform is transmitted is assumed to corrupt the signal with Gaussian noise. When the corrupted signal is received, it is first processed by the demodulator and then by the decoder.

Consider that a binary signal transmitted over a symbol interval $(0, T)$ is represented by $s_1(t)$ for a binary one and $s_2(t)$ for a binary zero. The received signal is $r(t) = s_i(t) + n(t)$, where $n(t)$ is a zero-mean Gaussian noise process. In Chapter 3 we described the detection of $r(t)$ in terms of two basic steps. In the first step, the received waveform is reduced to a single number, $z(T) = a_i + n_0$, where a_i is the signal component of $z(T)$ and n_0 is the noise component. The noise component, n_0, is a zero-mean *Gaussian random variable,* and thus $z(T)$ is a *Gaussian random variable* with a mean of either a_1 or a_2 depending on whether a binary one or binary zero was sent. In the second step of the detection process a decision was made as to which signal was transmitted, on the basis of comparing $z(T)$ to a threshold. The conditional probabilities of $z(T)$, $p(z|s_1)$, and $p(z|s_2)$ are shown in Figure 7.8, labeled likelihood of s_1 and likelihood of s_2. The demodulator in Figure 7.1, converts the set of time-ordered random variables $\{z(T)\}$ into a code sequence $\mathbf{Z}$, and passes it on to the decoder. The demodulator output can be configured in a variety of ways. It can be implemented to make a *firm or hard decision* as to whether $z(T)$ represents a zero or a one. In this case, the output of the demodulator is quantized to two levels, zero and one, and fed into the decoder (this is exactly the same threshold decision that was made in Chapters 3

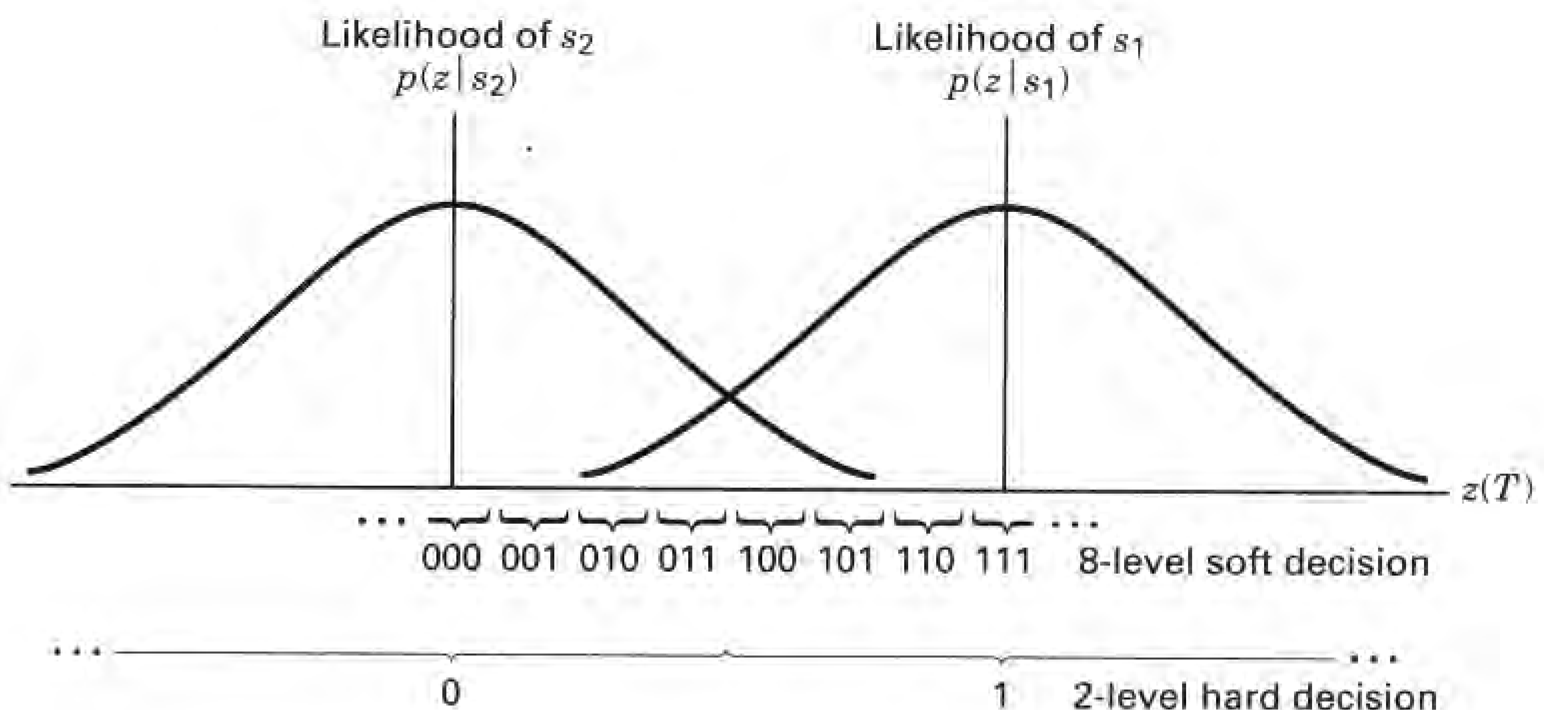

Figure 7.8 Hard and soft decoding decisions.

and 4). Since the decoder operates on the hard decisions made by the demodulator, the decoding is called *hard-decision decoding*.

The demodulator can also be configured to feed the decoder with a *quantized value* of $z(T)$ *greater than two levels*. Such an implementation furnishes the decoder with more information than is provided in the hard-decision case. When the quantization level of the demodulator output is greater than two, the decoding is called *soft-decision decoding*. Eight levels (3-bits) of quantization are illustrated on the abscissa of Figure 7.8. When the demodulator sends a hard binary decision to the decoder, it sends it a single binary symbol. When the demodulator sends a soft binary decision, quantized to eight levels, it sends the decoder a 3-bit word describing an interval along $z(T)$. In effect, sending such a 3-bit word in place of a single binary symbol is equivalent to sending the decoder a *measure of confidence* along with the code-symbol decision. Referring to Figure 7.8, if the demodulator sends 1 1 1 to the decoder, this is tantamount to declaring the code symbol to be a one with very high confidence, while sending a 1 0 0 is tantamount to declaring the code symbol to be a one with very low confidence. It should be clear that ultimately, every message decision out of the decoder must be a hard decision; otherwise, one might see computer printouts that read: "think it's a 1," "think it's a 0," and so on. The idea behind the demodulator *not making hard decisions* and sending more data (soft decisions) to the decoder can be thought of as an interim step to provide the decoder with more information, which the decoder then uses for recovering the message sequence (with better error performance than it could in the case of hard-decision decoding). In Figure 7.8, the 8-level soft-decision metric is often shown as $-7, -5, -3, -1, 1, 3, 5, 7$. Such a designation lends itself to a simple interpretation of the soft decision: The sign of the metric represents a decision (e.g., choose s_1 if positive, choose s_2 if negative), and the magnitude of the metric represents the confidence level of that decision. The only advantage for the metric shown in Figure 7.8 is that it avoids the use of negative numbers.

For a Gaussian channel, eight-level quantization results in a performance improvement of approximately 2 dB in required signal-to-noise ratio compared to two-level quantization. This means that eight-level soft-decision decoding can provide the same probability of bit error as that of hard-decision decoding, but requires 2 dB *less* E_b/N_0 for the same performance. Analog (or infinite-level quantization) results in a 2.2-dB performance improvement over two-level quantization; therefore, *eight-level quantization* results in a loss of approximately 0.2 dB compared to infinitely fine quantization. For this reason, quantization to more than eight levels can yield little performance improvement [12]. What price is paid for such improved soft-decision-decoder performance? In the case of hard-decision decoding, a single bit is used to describe each code symbol, while for eight-level quantized soft-decision decoding 3 bits are used to describe each code symbol; therefore, three times the amount of data must be handled during the decoding process. Hence the price paid for soft-decision decoding is an increase in required memory size at the decoder (and possibly a speed penalty).

Block decoding algorithms and convolutional decoding algorithms have been devised to operate with hard *or* soft decisions. However, soft-decision decoding is generally not used with block codes because it is considerably more difficult than

hard-decision decoding to implement. The most prevalent use of soft-decision decoding is with the *Viterbi convolutional decoding algorithm,* since with Viterbi decoding, soft decisions represent only a trivial increase in computation.

7.3.2.1 Binary Symmetric Channel

A binary symmetric channel (BSC) is a discrete memoryless channel (see Section 6.3.1) that has binary input and output alphabets and symmetric transition probabilities. It can be described by the conditional probabilities

$$\begin{aligned} P(0|1) &= P(1|0) = p \\ P(1|1) &= P(0|0) = 1 - p \end{aligned} \tag{7.5}$$

as illustrated in Figure 7.9. The probability that an output symbol will differ from the input symbol is p, and the probability that the output symbol will be identical to the input symbol is $(1 - p)$. The BSC is an example of a *hard-decision channel,* which means that, even though continuous-valued signals may be received by the demodulator, a BSC allows only firm decisions such that each demodulator output symbol, z_{ji}, as shown in Figure 7.1, consists of one of two binary values. The indexing of z_{ji} pertains to the jth code symbol of the ith branch word, Z_i. The demodulator then feeds the sequence $\mathbf{Z} = \{Z_i\}$ to the decoder.

Let $\mathbf{U}^{(m)}$ be a transmitted codeword over a BSC with symbol error probability p, and let $\mathbf{Z}$ be the corresponding received decoder sequence. As noted previously, a maximum likelihood decoder chooses the codeword $\mathbf{U}^{(m')}$ that maximizes the likelihood $P(\mathbf{Z}|\mathbf{U}^{(m)})$ or its logarithm. For a BSC, this is equivalent to choosing the codeword $\mathbf{U}^{(m')}$ that is closest in *Hamming distance* to $\mathbf{Z}$ [8]. Thus Hamming distance is an appropriate metric to describe the distance or closeness of fit between $\mathbf{U}^{(m)}$ and $\mathbf{Z}$. From all the possible transmitted sequences $\mathbf{U}^{(m)}$, the decoder chooses the $\mathbf{U}^{(m')}$ sequence for which the distance to $\mathbf{Z}$ is minimum.

Suppose that $\mathbf{U}^{(m)}$ and $\mathbf{Z}$ are each L-bit-long sequences and that they differ in d_m positions [i.e., the Hamming distance between $\mathbf{U}^{(m)}$ and $\mathbf{Z}$ is d_m]. Then, since the

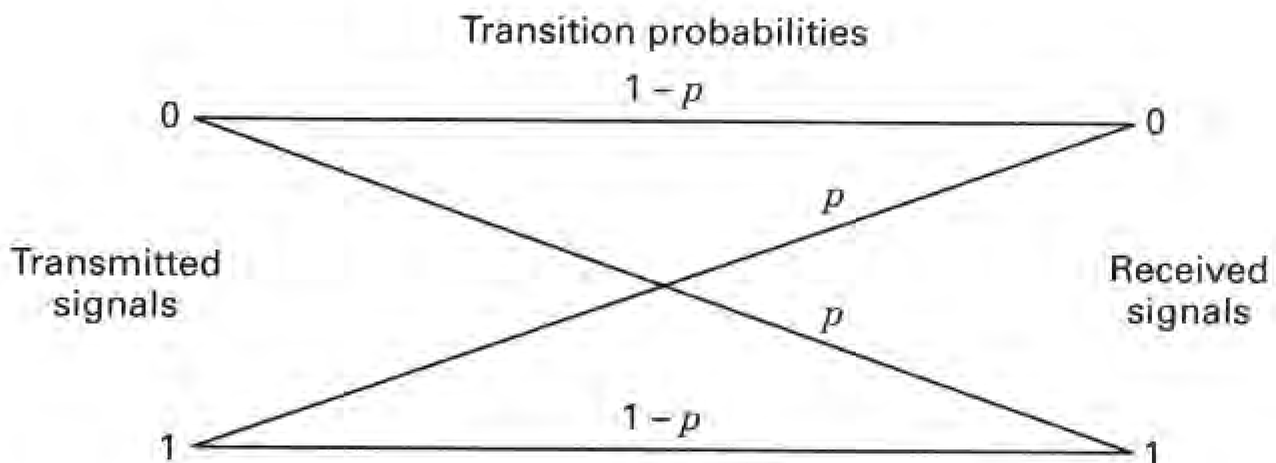

Figure 7.9 Binary symmetric channel (hard-decision channel).

channel is assumed to be memoryless, the probability that this $\mathbf{U}^{(m)}$ was transformed to the specific received $\mathbf{Z}$ at distance d_m from it can be written as

$$P(\mathbf{Z}|\mathbf{U}^{(m)}) = p^{d_m}(1-p)^{L-d_m} \tag{7.6}$$

and the log-likelihood function is

$$\log P(\mathbf{Z}|\mathbf{U}^{(m)}) = -d_m \log\left(\frac{1-p}{p}\right) + L\log(1-p) \tag{7.7}$$

If we compute this quantity for each possible transmitted sequence, the last term in the equation will be constant in each case. Assuming that $p < 0.5$, we can express Equation (7.7) as

$$\log P(\mathbf{Z}|\mathbf{U}^{(m)}) = -Ad_m - B \tag{7.8}$$

where A and B are positive constants. Therefore, choosing the codeword $\mathbf{U}^{(m')}$, such that the Hamming distance d_m to the received sequence $\mathbf{Z}$ is minimized, corresponds to *maximizing the likelihood or log-likelihood metric.* Consequently, over a BSC, the log-likelihood metric is conveniently replaced by the Hamming distance, and a maximum likelihood decoder will choose, in the tree or trellis diagram, the path whose corresponding sequence $\mathbf{U}^{(m')}$ is at the *minimum Hamming distance* to the received sequence $\mathbf{Z}$.

7.3.2.2 Gaussian Channel

For a Gaussian channel, each demodulator output symbol z_{ji}, as shown in Figure 7.1, is a value from a continuous alphabet. The symbol z_{ji} cannot be labeled as a correct or incorrect detection decision. Sending the decoder such soft decisions can be viewed as sending a family of conditional probabilities of the different symbols (see Section 6.3.1). It can be shown [8] that maximizing $P(\mathbf{Z}|\mathbf{U}^{(m)})$ is equivalent to maximizing the inner product between the codeword sequence $\mathbf{U}^{(m)}$ (consisting of binary symbols represented as bipolar values) and the analog-valued received sequence $\mathbf{Z}$. Thus, the decoder chooses the codeword $\mathbf{U}^{(m')}$ if it maximizes

$$\sum_{i=1}^{\infty} \sum_{j=1}^{n} z_{ji} u_{ji}^{(m)} \tag{7.9}$$

This is equivalent to choosing the codeword $\mathbf{U}^{(m')}$ that is closest in *Euclidean distance* to $\mathbf{Z}$. Even though the hard- and soft-decision channels require different metrics, the concept of choosing the codeword $\mathbf{U}^{(m')}$ that is closest to the received sequence, $\mathbf{Z}$, is the same in both cases. To implement the maximization of Equation (7.9) exactly, the decoder would have to be able to handle analog-valued arithmetic operations. This is impractical because the decoder is generally implemented digitally. Thus it is necessary to quantize the received symbols z_{ji}. Does Equation (7.9) remind you of the demodulation treatment in Chapters 3 and 4? Equation (7.9) is the discrete version of correlating an input received waveform, $r(t)$, with a reference waveform, $s_i(t)$, as expressed in Equation (4.15). The quantized Gaussian

channel, typically referred to as a *soft-decision channel,* is the channel model assumed for the soft-decision decoding described earlier.

7.3.3 The Viterbi Convolutional Decoding Algorithm

The Viterbi decoding algorithm was discovered and analyzed by Viterbi [13] in 1967. The Viterbi algorithm essentially performs maximum likelihood decoding; however, it reduces the computational load by taking advantage of the special structure in the code trellis. The advantage of Viterbi decoding, compared with brute-force decoding, is that the complexity of a Viterbi decoder is not a function of the number of symbols in the codeword sequence. The algorithm involves calculating a *measure of similarity, or distance,* between the received signal, at time t_i, and all the trellis paths entering each state at time t_i. The Viterbi algorithm removes from consideration those trellis paths that could not possibly be candidates for the maximum likelihood choice. When two paths enter the same state, the one having the best metric is chosen; this path is called the *surviving path.* This selection of surviving paths is performed for all the states. The decoder continues in this way to advance deeper into the trellis, making decisions by eliminating the least likely paths. The early rejection of the unlikely paths reduces the decoding complexity. In 1969, Omura [14] demonstrated that the Viterbi algorithm is, in fact, maximum likelihood. Note that the goal of selecting the optimum path can be expressed, equivalently, as choosing the codeword with the *maximum likelihood metric,* or as choosing the codeword with the *minimum distance metric.*

7.3.4 An Example of Viterbi Convolutional Decoding

For simplicity, a BSC is assumed; thus Hamming distance is a proper distance measure. The encoder for this example is shown in Figure 7.3, and the encoder trellis diagram is shown in Figure 7.7. A similar trellis can be used to represent the decoder, as shown in Figure 7.10. We start at time t_1 in the 00 state (flushing the encoder between messages provides the decoder with starting-state knowledge). Since in this example, there are only two possible transitions leaving any state, not all branches need be shown initially. The full trellis structure evolves after time t_3. The basic idea behind the decoding procedure can best be understood by examining the Figure 7.7 encoder trellis in concert with the Figure 7.10 decoder trellis. For the decoder trellis it is convenient at each time interval, to label each branch with the *Hamming distance* between the received code symbols and the branch word corresponding to the same branch from the encoder trellis. The example in Figure 7.10 shows a message sequence **m**, the corresponding codeword sequence **U**, and a noise corrupted received sequence **Z** = 11 01 01 10 01 . . . . The branch words seen on the *encoder trellis* branches characterize the encoder in Figure 7.3, and are known a priori to both the encoder and the decoder. These encoder branch words are the code symbols that would be expected to come from the encoder output as a result of each of the state transitions. The labels on the *decoder trellis* branches are accumulated by the decoder *on the fly.* That is, as the code symbols are received,

Input data sequence	**m**:	1	1	0	1	1	...
Transmitted codeword	**U**:	11	01	01	00	01	...
Received sequence	**Z**:	11	01	01	10	01	...

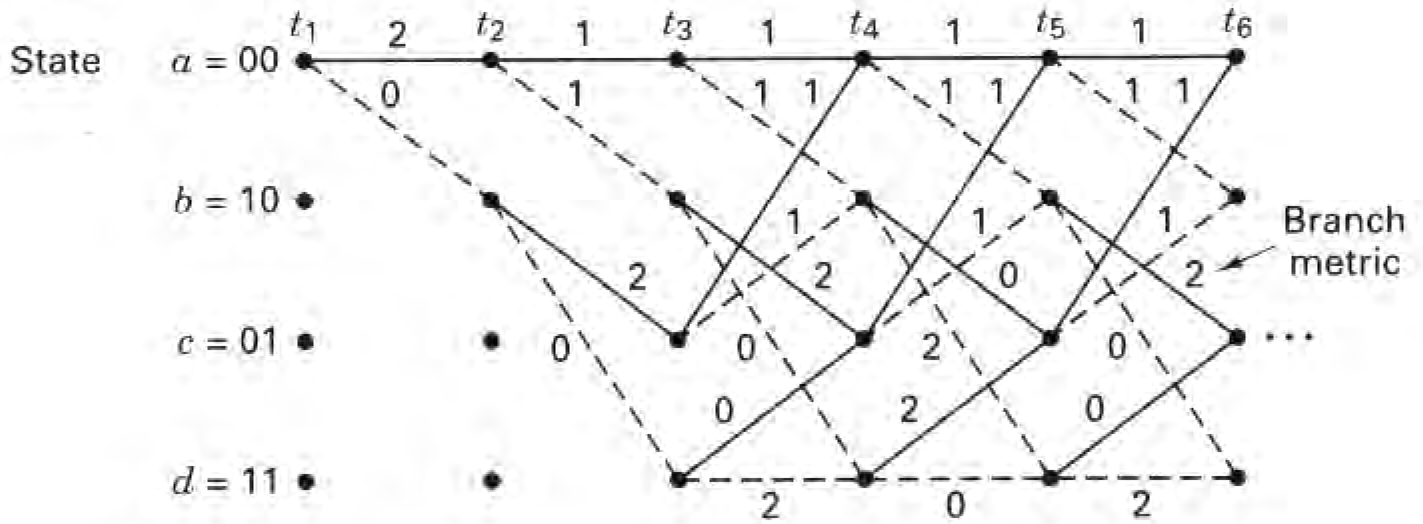

Figure 7.10 Decoder trellis diagram (rate $\frac{1}{2}$, $K = 3$).

each branch of the decoder trellis is labeled with a metric of similarity (Hamming distance) between the received code symbols and each of the branch words for that time interval. From the received sequence **Z**, shown in Figure 7.10, we see that the code symbols received at (following) time t_1 are 11. In order to label the decoder branches at (departing) time t_1 with the appropriate Hamming distance metric, we look at the Figure 7.7 encoder trellis. Here we see that a state $00 \rightarrow 00$ transition yields an output branch word of 00. But we received 11. Therefore, on the decoder trellis we label the state $00 \rightarrow 00$ transition with Hamming distance between them, namely 2. Looking at the encoder trellis again, we see that a state $00 \rightarrow 10$ transition yields an output branch word of 11, which corresponds exactly with the code symbols we received at time t_1. Therefore, on the decoder trellis, we label the state $00 \rightarrow 10$ transition with a Hamming distance of 0. In summary, the metric entered on a decoder trellis branch represents the difference (distance) between what was received and what "should have been" received had the branch word associated with that branch been transmitted. In effect, these metrics describe a correlation-like measure between a received branch word and each of the candidate branch words. We continue labeling the decoder trellis branches in this way as the symbols are received at each time t_i. The decoding algorithm uses these Hamming distance metrics to find the *most likely* (minimum distance) path through the trellis.

The basis of *Viterbi decoding* is the following observation: If any two paths in the trellis merge to a single state, one of them can always be eliminated in the search for an optimum path. For example, Figure 7.11 shows two paths merging at time t_5 to state 00. Let us define the *cumulative Hamming path metric* of a given path at time t_i as the sum of the branch Hamming distance metrics along that path up to time t_i. In Figure 7.11 the upper path has metric 4; the lower has metric 1. The upper path cannot be a portion of the optimum path because the lower path, which

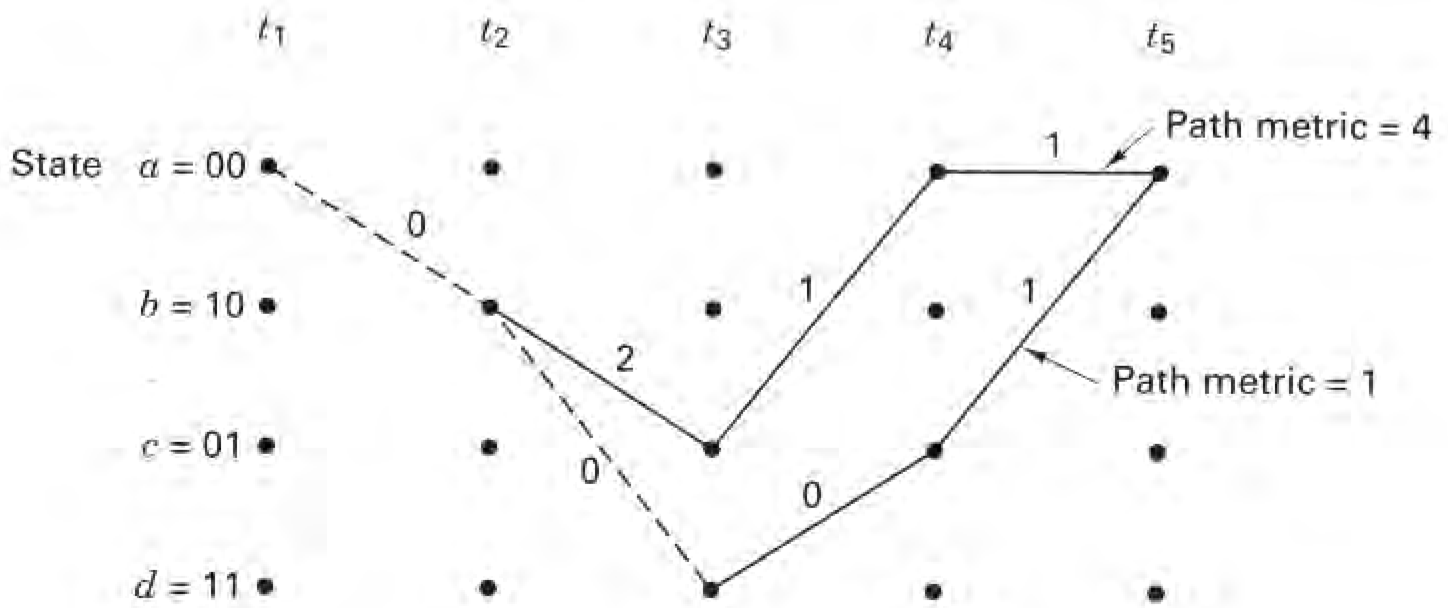

Figure 7.11 Path metrics for two merging paths.

enters the same state, has a lower metric. This observation holds because of the Markov nature of the encoder state: The present state summarizes the encoder history in the sense that previous states cannot affect future states or future output branches.

At each time t_i there are 2^{K-1} states in the trellis, where K is the constraint length, and each state can be *entered by means of two paths.* Viterbi decoding consists of computing the metrics for the two paths entering each state and *eliminating one of them.* This computation is done for each of the 2^{K-1} states or nodes at time t_i; then the decoder moves to time t_{i+1} and repeats the process. At a given time, the winning path metric for each state is designated as the *state metric* for that state at that time. The first few steps in our decoding example are as follows (see Figure 7.12). Assume that the input data sequence **m**, codeword **U**, and received sequence **Z** are as shown in Figure 7.10. Assume that the decoder knows the correct initial state of the trellis. (This assumption is not necessary in practice, but simplifies the explanation.) At time t_1 the received code symbols are 11. From state 00 the only possible transitions are to state 00 or state 10, as shown in Figure 7.12a. State 00 $\rightarrow$ 00 transition has branch metric 2; state 00 $\rightarrow$ 10 transition has branch metric 0. At time t_2 there are two possible branches leaving each state, as shown in Figure 7.12b. The cumulative metrics of these branches are labeled state metrics Γ_a, Γ_b, Γ_c, and Γ_d, corresponding to the terminating state. At time t_3 in Figure 7.12c there are again two branches diverging from each state. As a result, there are two paths entering each state at time t_4. One path entering each state can be eliminated, namely, the one having the larger cumulative path metric. Should metrics of the two entering paths be of equal value, one path is chosen for elimination by using an arbitrary rule. The surviving path into each state is shown in Figure 7.12d. At this point in the decoding process, there is only a single surviving path, termed the *common stem,* between times t_1 and t_2. Therefore, the decoder can now decide that the state transition which occurred between t_1 and t_2 was 00 $\rightarrow$ 10. Since this transition is produced by an input bit one, the decoder outputs a one as the first decoded bit.

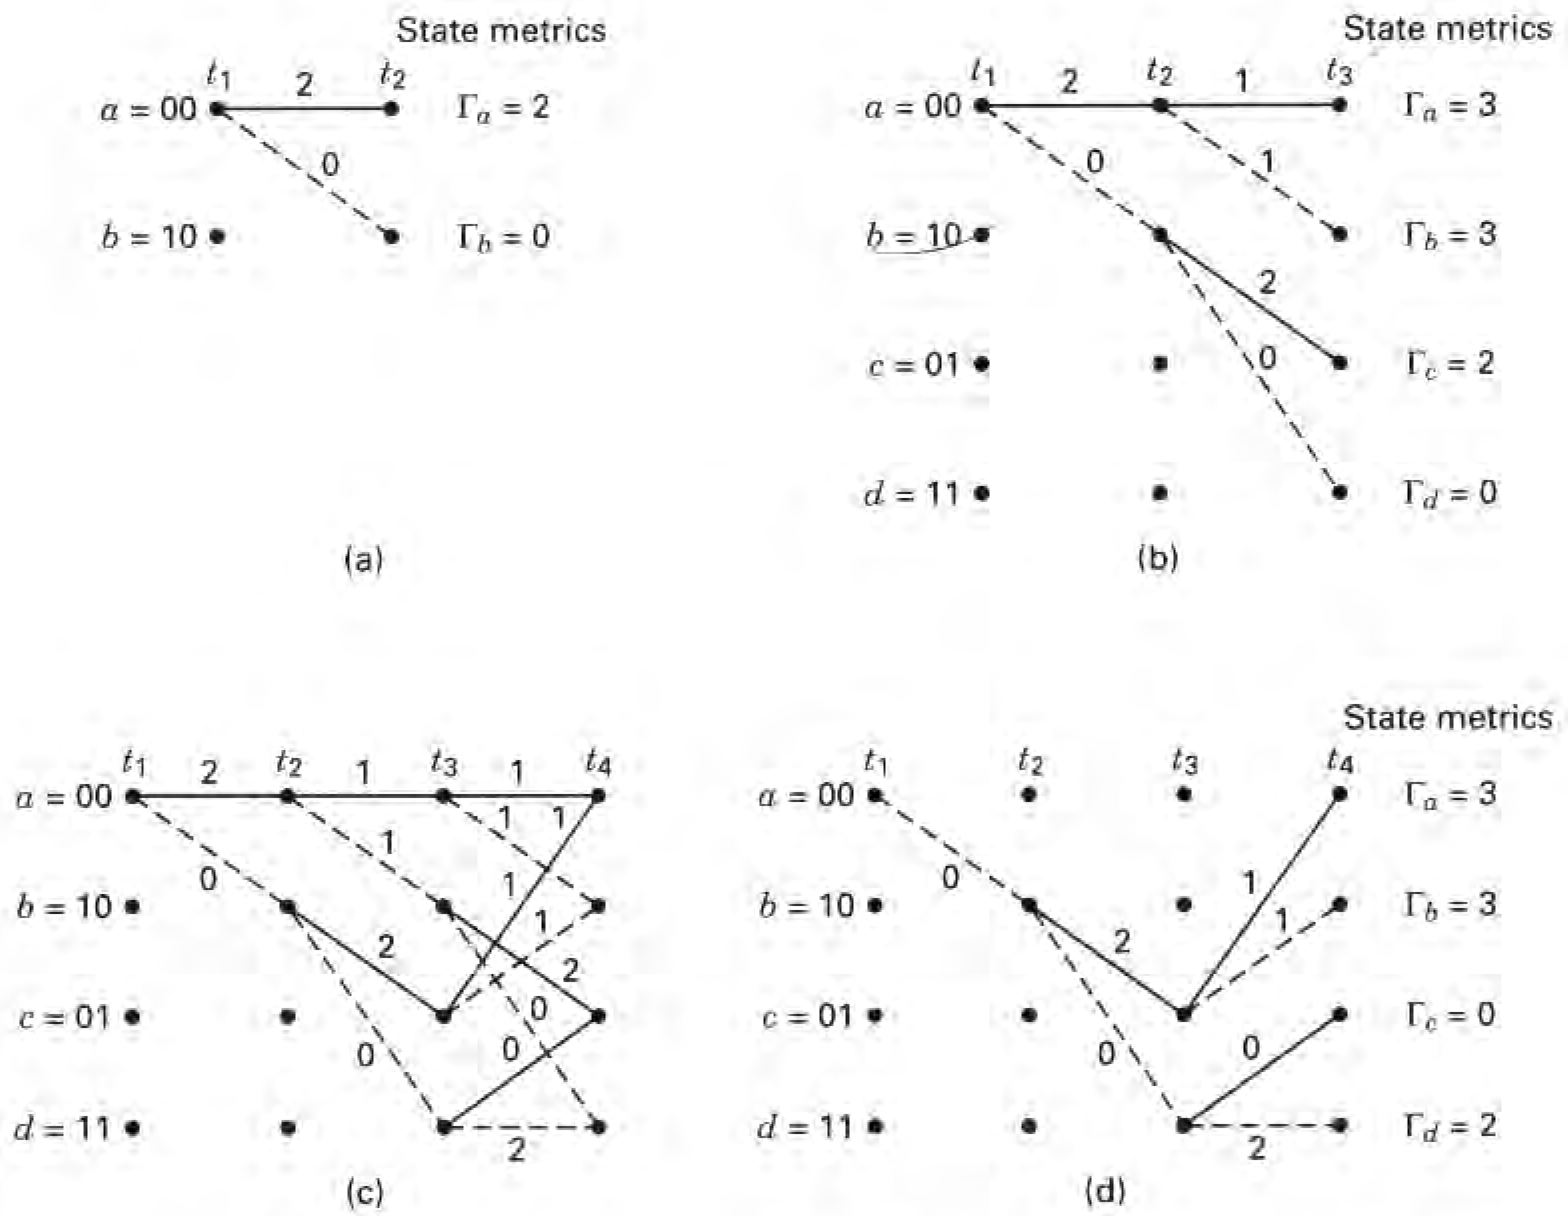

Figure 7.12 Selection of survivor paths, (a) Survivors at t_2. (b) Survivors at t_3. (c) Metric comparisons at t_4. (d) Survivors at t_4. (e) Metric comparisons at t_5. (f) Survivors at t_5. (g) Metric comparisons at t_6. (h) Survivors at t_6.

Here we can see how the decoding of the surviving branch is facilitated by having drawn the trellis branches with solid lines for input zeros and dashed lines for input ones. Note that the first bit was not decoded until the path metric computation had proceeded to a much greater depth into the trellis. For a typical decoder implementation, this represents a decoding delay which can be as much as five times the constraint length in bits.

At each succeeding step in the decoding process, there will always be two possible paths entering each state; one of the two will be eliminated by comparing the path metrics. Figure 7.12e shows the next step in the decoding process. Again, at time t_5 there are two paths entering each state, and one of each pair can be eliminated. Figure 7.12f shows the survivors at time t_5. Notice that in our example we cannot yet make a decision on the second input data bit because there still are two paths leaving the state 10 node at time t_2. At time t_6 in Figure 7.12g we again see the pattern of remerging paths, and in Figure 7.12h we see the survivors at time t_6. Also, in Figure 7.12h the decoder outputs one as the second decoded bit, corre-

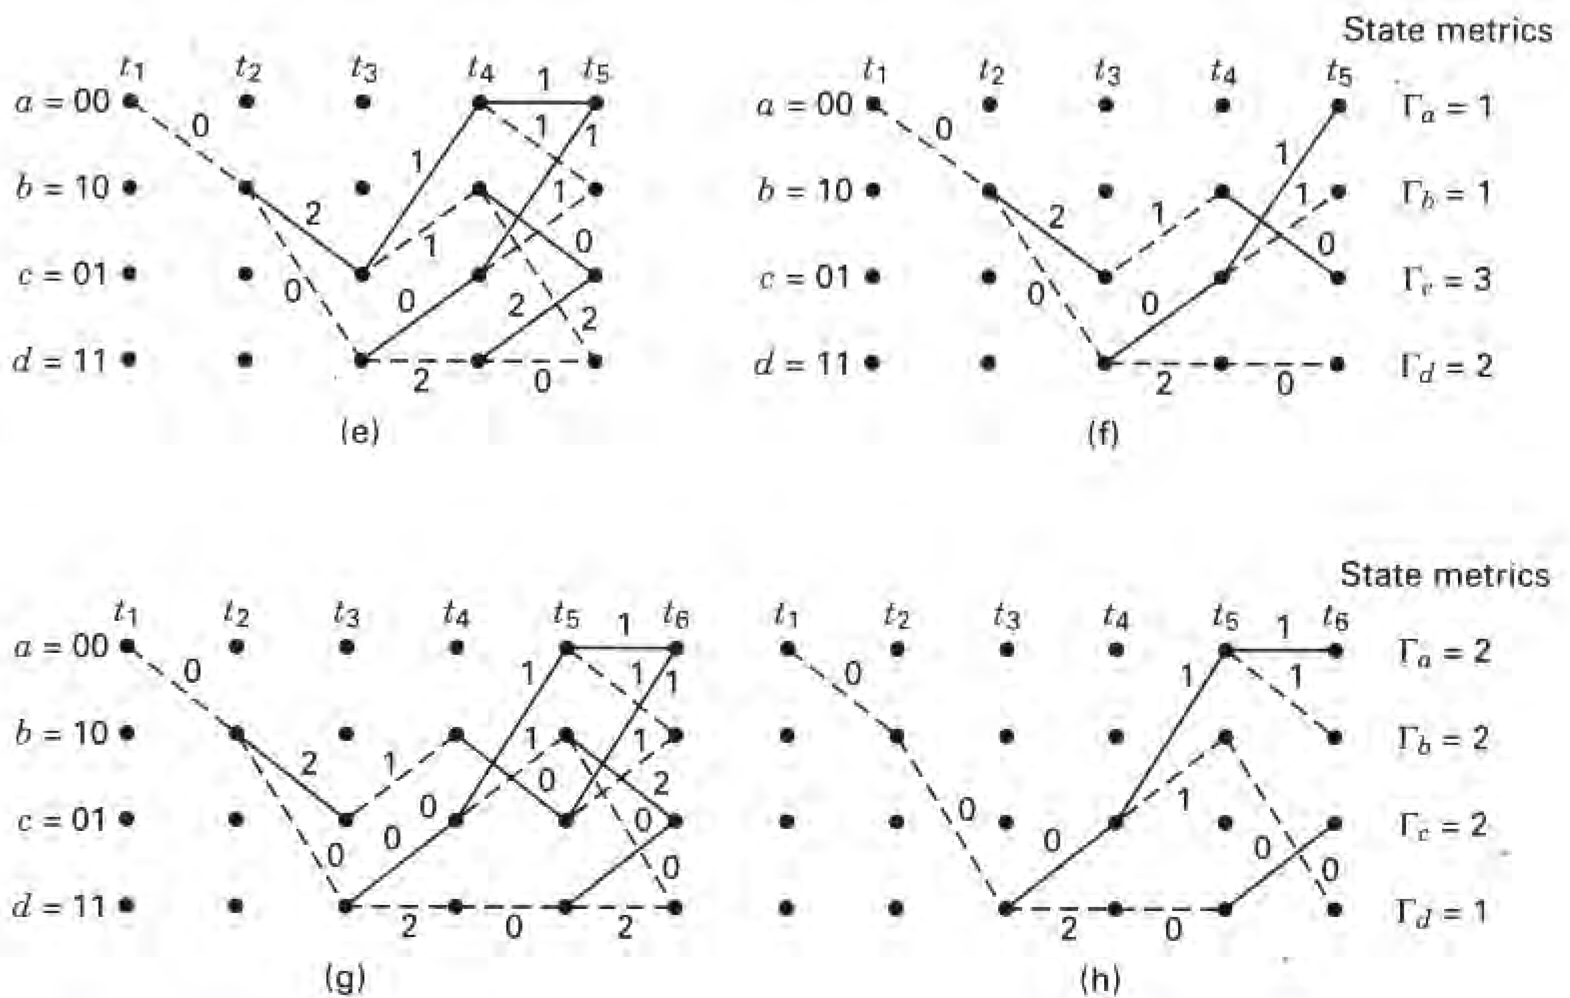

Figure 7.12 (*Continued*)

sponding to the single surviving path between t_2 and t_3. The decoder continues in this way to advance deeper into the trellis and to make decisions on the input data bits by eliminating all paths but one.

Pruning the trellis (as paths remerge) guarantees that there are never more paths than there are states. For this example, verify that after each pruning in Figures 7.12b, d, f, and h, there are only 4 paths. Compare this to attempting a "brute force" maximum-likelihood sequence estimation without using the Viterbi algorithm. In that case, the number of possible paths (representing possible sequences) is an exponential function of sequence length. For a binary codeword sequence that has a length of L branch words, there are 2^L possible sequences.

7.3.5 Decoder Implementation

In the context of the trellis diagram of Figure 7.10, transitions during any one time interval can be grouped into $2^{\nu-1}$ disjoint cells, each cell depicting four possible transitions, where $\nu = K - 1$ is called the *encoder memory*. For the $K = 3$ example, $\nu = 2$ and $2^{\nu-1} = 2$ cells. These cells are shown in Figure 7.13, where a, b, c, and d refer to the states at time t_i, and a', b', c', and d' refer to the states at time t_{i+1}. Shown on each transition is the *branch metric* δ_{xy}, where the subscript indicates that the metric corresponds to the transition from state x to state y. These cells and the associated logic units that update the *state metrics* $\{\Gamma_x\}$, where x designates a particular state, represent the basic building blocks of the decoder.

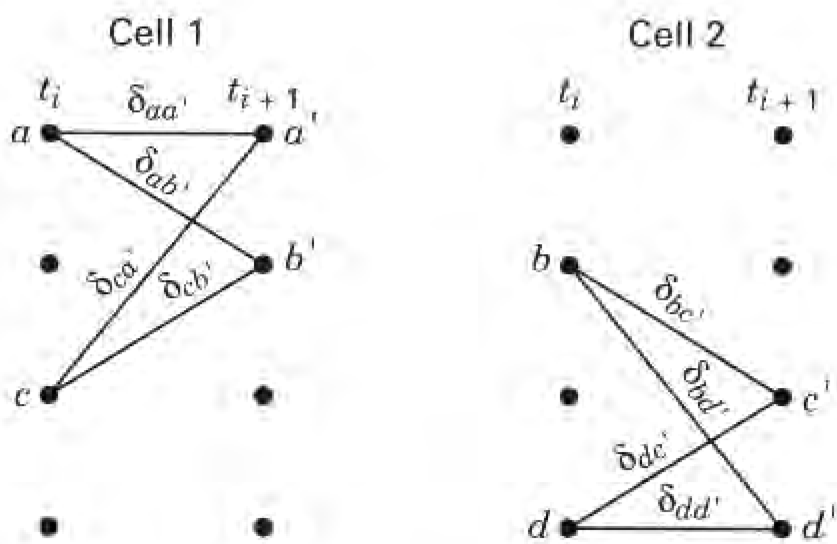

Figure 7.13 Example of decoder cells.

7.3.5.1 Add-Compare-Select Computation

Continuing with the $K = 3$, 2-cell example, Figure 7.14 illustrates the logic unit that corresponds to cell 1. The logic executes the special purpose computation called *add-compare-select* (ACS). The state metric $\Gamma_{a'}$ is calculated by adding the previous-time state metric of state a, Γ_a, to the branch metric $\delta_{aa'}$ and the previous-time state metric of state c, Γ_c, to the branch metric $\delta_{ca'}$. This results in two possible path metrics as candidates for the new state metric $\Gamma_{a'}$. The two candidates are compared in the logic unit of Figure 7.14. The largest likelihood (smallest distance) of the two path metrics is stored as the new state metric $\Gamma_{a'}$ for state a. Also stored is the *new* path history $\hat{m}_{a'}$ for state a, where $\hat{m}_{a'}$ is the message-path history of the state augmented by the data of the winning path.

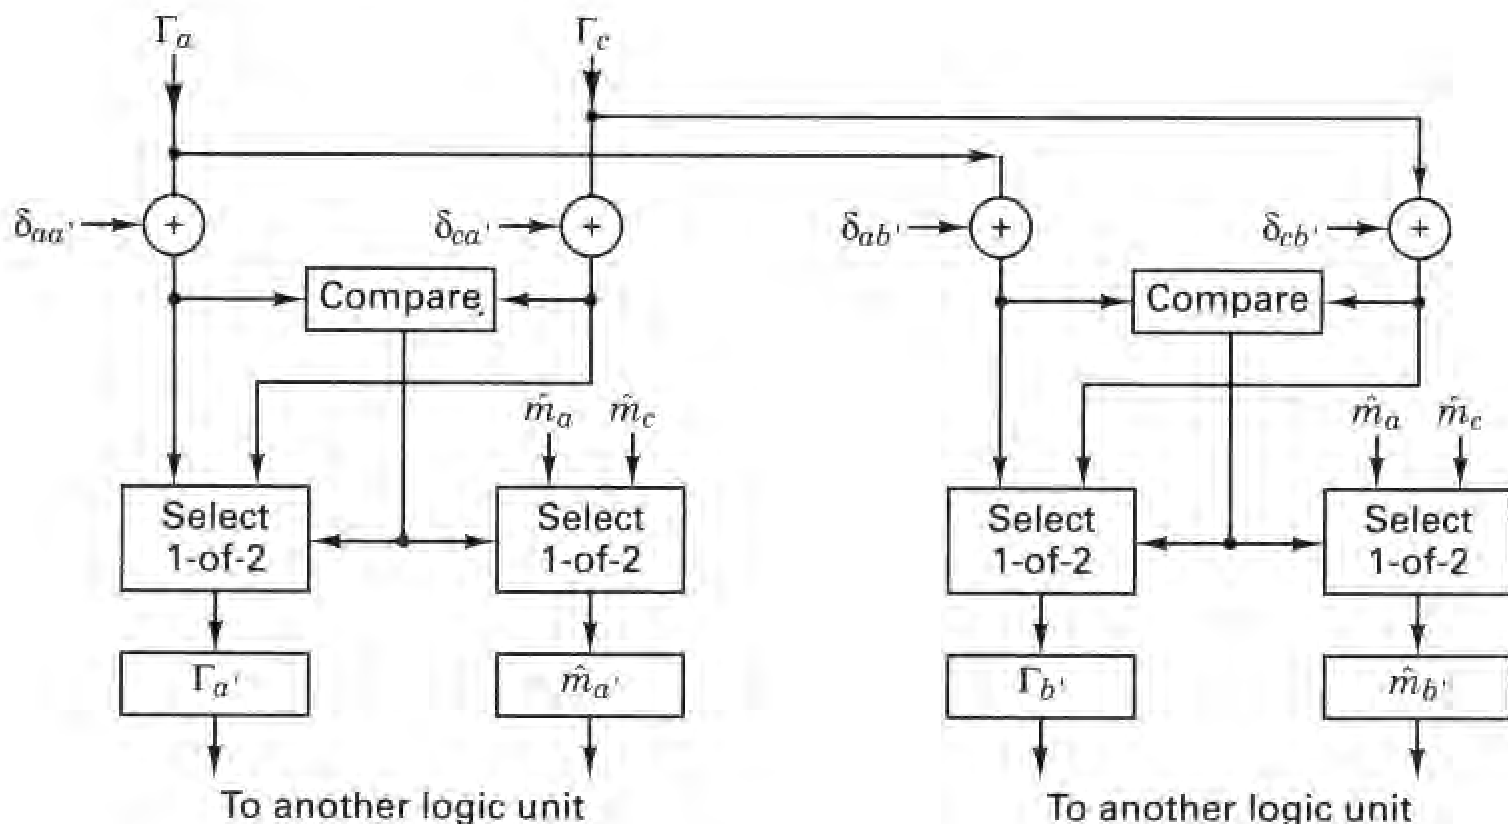

Figure 7.14 Logic unit that implements the add-compare-select functions corresponding to cell #1.

Also shown in Figure 7.14 is the cell-1 ACS logic that yields the new state metric $\Gamma_{b'}$ and the new path history $\hat{m}_{b'}$. This ACS operation is similarly performed for the paths in other cells. The oldest bit on the path with the smallest state metric forms the decoder output.

7.3.5.2 Add-Compare-Select as seen on the Trellis

Consider the same example that was used for describing Viterbi decoding in Section 7.3.4. The message sequence was **m** = 1 1 0 1 1, the codeword sequence was **U** = 11 01 01 00 01, and the received sequence was **Z** = 11 01 01 10 01. Figure 7.15 depicts a decoding trellis diagram similar to Figure 7.10. A branch metric that labels each branch is the Hamming distance between the received code symbols and the corresponding branch word from the encoder trellis. Additionally, the Figure 7.15 trellis indicates a value at each state x, and for each time from time t_2 to t_6, which is a state metric Γ_x. We perform the add-compare-select (ACS) operation when there are two transitions entering a state, as there are for times t_4 and later. For example at time t_4, the value of the state metric for state a is obtained by incrementing the state metric $\Gamma_a = 3$ at time t_3 with the branch metric $\delta_{aa'} = 1$ yielding a candidate value of 4. Simultaneously, the state metric $\Gamma_c = 2$ at time t_3 is incremented with the branch metric $\delta_{ca'} = 1$ yielding a candidate value of 3. The select operation of the ACS process selects the largest-likelihood (minimum distance) path metric as the new state metric; hence, for state a at time t_4, the new state metric is $\Gamma_{a'} = 3$. The winning path is shown with a heavy line and the path that has been dropped is shown with a lighter line. On the trellis of Figure 7.15, observe the state metrics from left to right. Verify that at each time, the value of each state metric is obtained by incrementing the connected state metric from the previous time along the winning path (heavy line) with the branch metric between them. At some

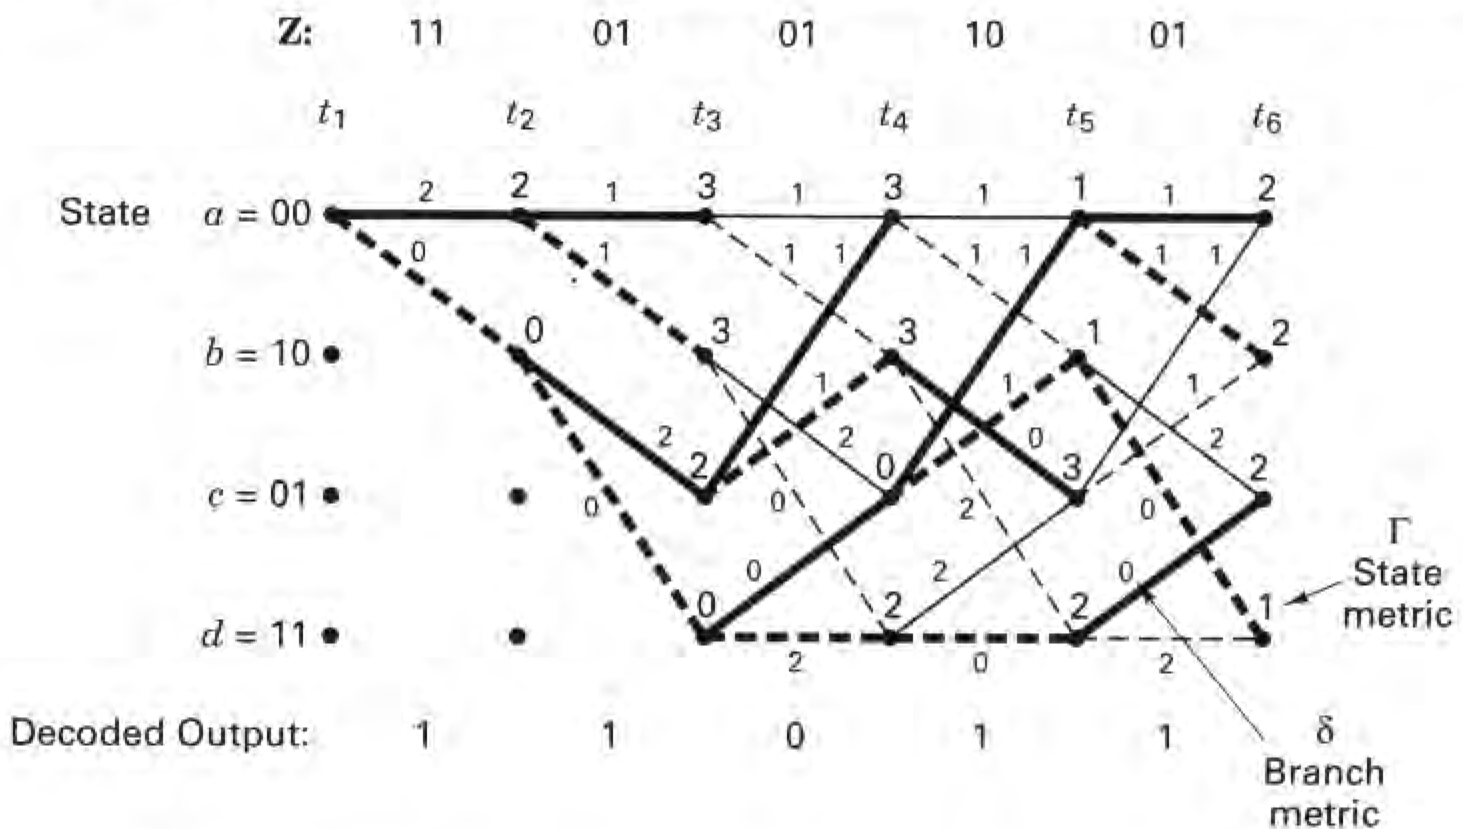

Figure 7.15 Add-compare-select computations in Viterbi decoding.

point in the trellis (after a time interval of 4 or 5 times the constraint length), the oldest bits can be decoded. As an example, looking at time t_6 in Figure 7.15, we see that the minimum-distance state metric has a value of 1. From this state d, the winning path can be traced back to time t_1, and one can verify that the decoded message is the same as the original message, by the convention that dashed and solid lines represent binary ones and zeros respectively.

7.3.6 Path Memory and Synchronization

The storage requirements of the Viterbi decoder grow exponentially with constraint length K. For a code with rate $1/n$, the decoder retains a set of 2^{K-1} paths after each decoding step. With high probability, these paths will not be mutually disjoint very far back from the present decoding depth [12]. All of the 2^{K-1} paths tend to have a common stem which eventually branches to the various states. Thus if the decoder stores enough of the history of the 2^{K-1} paths, the oldest bits on all paths will be the same. A simple decoder implementation, then, contains a *fixed amount of path history* and outputs the oldest bit on an arbitrary path each time it steps one level deeper into the trellis. The amount of path storage required is [12]

$$u = h2^{K-1} \tag{7.10}$$

where h is the length of the information bit path history per state. A refinement, which minimizes the value of h, uses the oldest bit on the most likely path as the decoder output, instead of the oldest bit on an arbitrary path. It has been demonstrated [12] that a value of h of 4 or 5 times the code constraint length is sufficient for near-optimum decoder performance. The storage requirement u is the basic limitation on the implementation of Viterbi decoders. Commercial decoders are limited to a constraint length of about $K = 10$. Efforts to increase coding gain by further increasing constraint length are met by the exponential increase in memory requirements (and complexity) that follows from Equation (7.10).

Branch word synchronization is the process of determining the beginning of a branch word in the received sequence. Such synchronization can take place without new information being added to the transmitted symbol stream because the received data appear to have an excessive error rate when not synchronized. Therefore, a simple way of accomplishing synchronization is to monitor some concomitant indication of this large error rate, that is, the rate at which the state metrics are increasing or the rate at which the surviving paths in the trellis merge. The monitored parameters are compared to a threshold, and synchronization is then adjusted accordingly.

7.4 PROPERTIES OF CONVOLUTIONAL CODES

7.4.1 Distance Properties of Convolutional Codes

Consider the distance properties of convolutional codes in the context of the simple encoder in Figure 7.3 and its trellis diagram in Figure 7.7. We want to evaluate the distance between all possible pairs of codeword sequences. As in the case of

block codes (see Section 6.5.2), we are interested in the *minimum distance* between all pairs of such codeword sequences in the code, since the minimum distance is related to the error-correcting capability of the code. Because a convolutional code is a group or *linear code* [6], there is no loss in generality in simply finding the minimum distance between each of the codeword sequences and the all-zeros sequence. In other words, for a linear code, any test message is just as "good" as any other test message. So, why not choose one that is easy to keep track of—namely, the all-zeros sequence? Assuming that the all-zeros input sequence was transmitted, the paths of interest are those that start and end in the 00 state and do not return to the 00 state anywhere in between. An error will occur whenever the distance of any other path that merges with the $a = 00$ state at time t_i is less than that of the all-zeros path up to time t_i, causing the all-zeros path to be discarded in the decoding process. In other words, given the all-zeros transmission, an error occurs whenever the *all-zeros path does not survive.* Thus, an error of interest is associated with a surviving path that diverges from and then remerges to the all-zeros path. One might ask, Why is it necessary for the path to remerge? Isn't the divergence enough to indicate an error? Yes, of course, but an error characterized by *only* a divergence means that the decoder, from that point on, will be outputting "garbage" for the rest of the message duration. We want to quantify the decoder's capability in terms of errors that will usually take place—that is, we want to learn the "easiest" way for the decoder to make an error. The minimum distance for making such an error can be found by exhaustively examining every path from the 00 state to the 00 state. First, let us redraw the trellis diagram, shown in Figure 7.16, labeling each branch with its Hamming distance from the all-zeros codeword instead of with its branch word symbols. The Hamming distance between two unequal-length sequences will be found by first appending the necessary number of zeros to the shorter sequence to make the two sequences equal in length. Consider all the paths that diverge from the all-zeros path and then remerge for the first time at some arbitrary node. From Figure 7.16 we can compute the distances of these paths from the all-zeros path. There is one path at distance 5 from the all-zeros path; this path departs from the all-zeros path at time t_1 and merges with it at time t_4. Similarly, there are two paths at distance 6, one which departs at time t_1 and merges at time t_5, and the other which departs at time t_1 and merges at time t_6, and so on. We can also see from the dashed and solid lines of the diagram that the input bits for the distance 5 path are 1 0 0; it differs in only one input bit from the all-zeros input sequence. Similarly, the input bits for the distance 6 paths are 1 1 0 0 and 1 0 1 0 0; each differs in two positions from the all-zeros path. The minimum distance in the set of all arbitrarily long paths that diverge and remerge, called the *minimum free distance,* or simply the *free distance,* is seen to be 5 in this example, as shown with the heavy lines in Figure 7.16. For calculating the error-correcting capability of the code, we repeat Equation (6.44) with the minimum distance $d_{\min}$ replaced by the free distance d_f as

$$t = \left\lfloor \frac{d_f - 1}{2} \right\rfloor \tag{7.11}$$

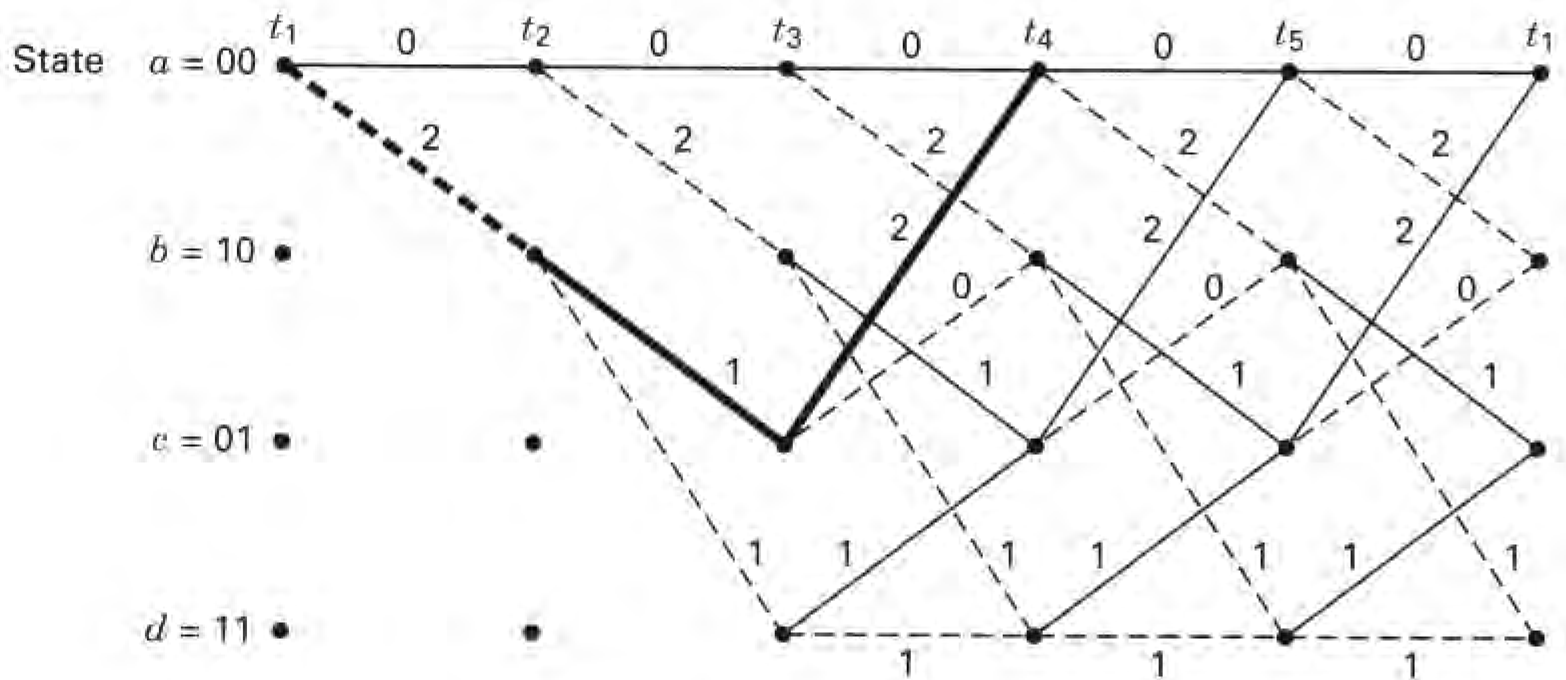

Figure 7.16 Trellis diagram, labeled with distances from the all-zeros path.

where $\lfloor x \rfloor$ means the largest integer no greater than x. Setting $d_f = 5$, we see that the code, characterized by the Figure 7.3 encoder, can correct any two channel errors. (See Section 7.4.1.1.)

A trellis diagram represents "the rules of the game." It is a shorthand description of all the possible transitions and their corresponding start and finish states associated with a particular finite-state machine. The trellis diagram offers some insight into the benefit (coding gain) when using error-correction coding. Consider Figure 7.16 and the possible divergence-remergence error paths. From this picture one sees that the decoder cannot make an error in any *arbitrary way*. The error path must follow one of the allowable transitions. The trellis pinpoints all such allowable paths. By having encoded the data in this way, we have placed constraints on the transmitted signal. The decoder knows these constraints, and this knowledge enables the system to more easily (using less E_b/N_0) meet some error performance requirements.

Although Figure 7.16 presents the computation of free distance in a straightforward way, a more direct closed-form expression can be obtained by starting with the state diagram in Figure 7.5. First, we label the branches of the state diagram as either $D^0 = 1$, D^1, or D^2, shown in Figure 7.17, where the exponent of D denotes the Hamming distance from the branch word of that branch to the all-zeros branch. The self-loop at node a can be eliminated since it contributes nothing to the distance properties of a codeword sequence relative to the all-zeros sequence. Furthermore, node a can be split into two nodes (labeled a and e), one of which represents the input and the other the output of the state diagram. All paths originating at $a = 00$ and terminating at $e = 00$ can be traced on the modified state diagram of Figure 7.17. We can calculate the transfer function of path $a\ b\ c\ e$ (starting and ending at state 00) in terms of the indeterminate "placeholder" D, as $D^2\ D\ D^2 = D^5$. The exponent of D represents the cumulative tally of the number of ones in the path, and hence the Hamming distance from the all-zeros path. Similarly, the

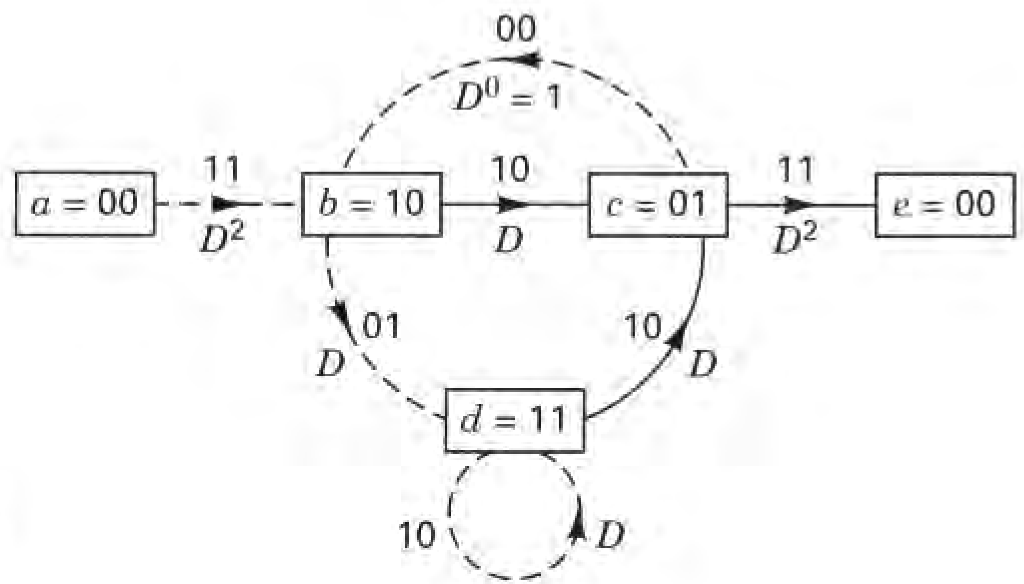

Figure 7.17 State diagram, labeled according to distance from the all-zeros path.

paths $a\ b\ d\ c\ e$ and $a\ b\ c\ b\ c\ e$ each have the transfer function D^6 and thus a Hamming distance of 6 from the all-zeros path. We now write the state equations as

$$\begin{aligned} X_b &= D^2X_a + X_c \\ X_c &= DX_b + DX_d \\ X_d &= DX_b + DX_d \\ X_e &= D^2X_c \end{aligned} \tag{7.12}$$

where $X_a, \ldots, X_e$ are dummy variables for the partial paths to the intermediate nodes. The *transfer function, T(D)*, sometimes called the *generating function* of the code can be expressed as $T(D) = X_e/X_a$. By solving the state equations shown in Equation (7.12), we obtain [15, 16]

$$\begin{aligned} T(D) &= \frac{D^5}{1 - 2D} \\ &= D^5 + 2D^6 + 4D^7 + \cdots + 2^\ell D^{\ell+5} + \cdots \end{aligned} \tag{7.13}$$

The transfer function for this code indicates that there is a single path of distance 5 from the all-zeros path, two of distance 6, four of distance 7, and in general, there are 2^ℓ paths of distance $\ell + 5$ from the all-zeros path, where $\ell = 0, 1, 2, \ldots$. The free distance d_f of the code is the Hamming weight of the lowest-order term in the expansion of $T(D)$. In this example $d_f = 5$. In evaluating distance properties, the transfer function, $T(D)$, cannot be used for long constraint lengths since the complexity of $T(D)$ increases exponentially with constraint length.

The transfer function can be used to provide more detailed information than just the distance of the various paths. Let us introduce a factor L into each branch of the state diagram so that the exponent of L can serve as a counter to indicate the number of branches in any given path from state $a = 00$ to state $e = 00$. Furthermore, we can introduce a factor N into all branch transitions caused by the input bit one. Thus, as each branch is traversed, the cumulative exponent on N increases by one, only if that branch transition is due to an input bit one. For the convolutional

code characterized in the Figure 7.3 example, the additional factors L and N are shown on the modified state diagram of Figure 7.18. Equations (7.12) can now be modified as follows:

$$\begin{aligned} X_b &= D^2LNX_a + LNX_c \\ X_c &= DLX_b + DLX_d \\ X_d &= DLNX_b + DLNX_d \\ X_e &= D^2LX_c \end{aligned} \tag{7.14}$$

The transfer function of this augmented state diagram is

$$\begin{aligned} T(D, L, N) &= \frac{D^5L^3N}{1 - DL(1 + L)N} \\ &= D^5L^3N + D^6L^4(1 + L)N^2 + D^7L^5(1 + L)^2N^3 \\ &\quad + \cdots + D^{\ell+5}L^{\ell+3}N^{\ell+1} + \cdots \end{aligned} \tag{7.15}$$

Thus, we can verify some of the path properties displayed in Figure 7.16. There is one path of distance 5, length 3, which differs in one input bit from the all-zeros path. There are two paths of distance 6, one of which is length 4, the other length 5, and both differ in two input bits from the all-zeros path. Also, of the distance 7 paths, one is of length 5, two are of length 6, and one is of length 7; all four paths correspond to input sequences that differ in three input bits from the all-zeros path. Thus if the all-zeros path is the correct path and the noise causes us to choose one of the incorrect paths of distance 7, three bit errors will be made.

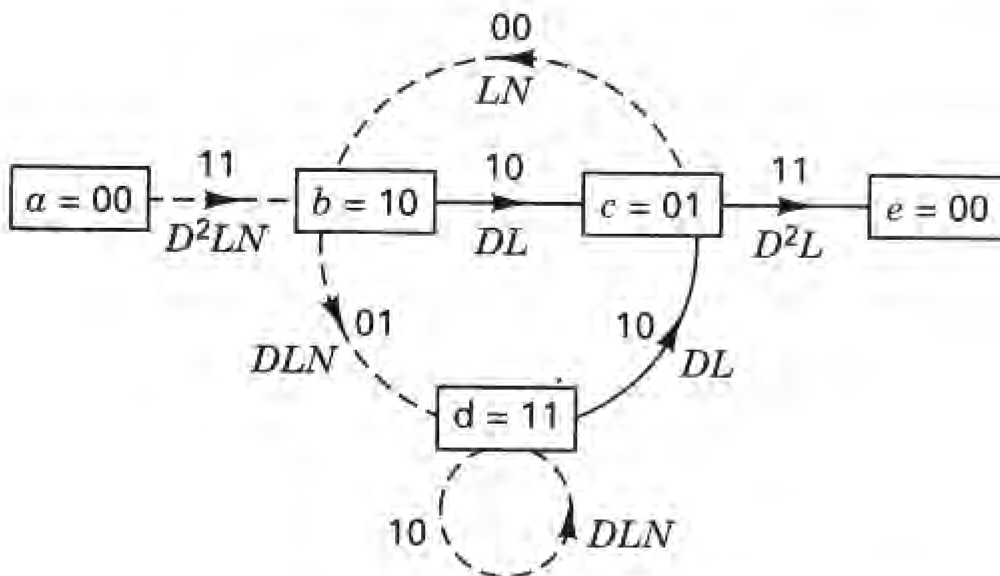

Figure 7.18 State diagram, labeled according to distance, length, and number of input ones.

7.4.1.1 Error-Correcting Capability of Convolutional Codes

In the study of block codes in Chapter 6, we saw that the error-correcting capability, t, represented the number of code symbol errors that could, with maximum likelihood decoding, be corrected in each block length of the code. However, when decoding convolutional codes, the error-correcting capability cannot be stated so succinctly. With regard to Equation (7.11), we can say that the code can, with maximum likelihood decoding, correct t errors within a few constraint lengths,

where "few" here means 3 to 5. The exact length depends on how the errors are distributed. For a particular code and error pattern, the length can be bounded using transfer function methods. Such bounds are described later.

7.4.2 Systematic and Nonsystematic Convolutional Codes

A *systematic* convolutional code is one in which the input k-tuple appears as part of the output branch word n-tuple associated with that k-tuple. Figure 7.19 shows a binary, rate $\frac{1}{2}$, $K = 3$ systematic encoder. For linear block codes, any nonsystematic code can be transformed into a systematic code with the same block distance properties. This is not the case for convolutional codes. The reason for this is that convolutional codes depend largely on *free distance*; making the convolutional code systematic, in general, *reduces* the maximum possible free distance for a given constraint length and rate.

Table 7.1 shows the maximum free distance for rate $\frac{1}{2}$ systematic and nonsystematic codes for $K = 2$ through 8. For large constraint lengths the results are even more widely separated [17].

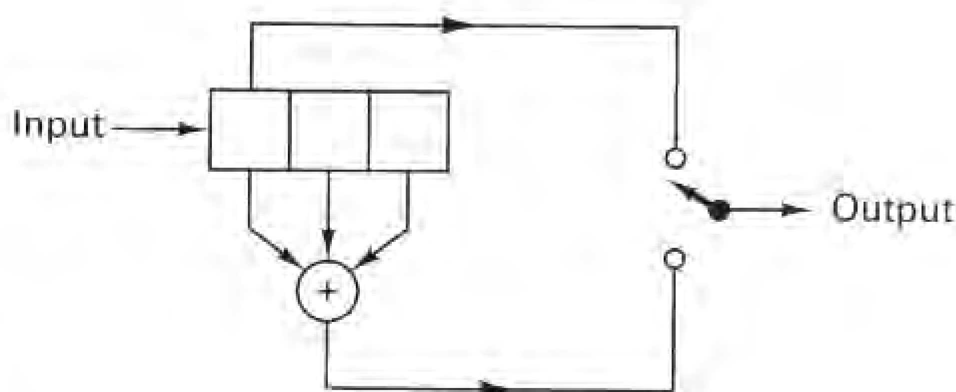

Figure 7.19 Systematic convolutional encoder, rate $\frac{1}{2}$, $K = 3$.

TABLE 7.1 Comparison of Systematic and Nonsystematic Free Distance, Rate $\frac{1}{2}$

Constraint Length	Free Distance Systematic	Free Distance Nonsystematic
2	3	3
3	4	5
4	4	6
5	5	7
6	6	8
7	6	10
8	7	10

Source: A. J. Viterbi and J. K. Omura, *Principles of Digital Communication and Coding*, McGraw-Hill Book Company, New York, 1979, p. 251.

7.4.3 Catastrophic Error Propagation in Convolutional Codes

A *catastrophic error* is defined as an event whereby a finite number of code symbol errors cause an infinite number of decoded data bit errors. Massey and Sain [18] have derived a necessary and sufficient condition for convolutional codes to display catastrophic error propagation. For rate $1/n$ codes with register taps designated by polynomial generators, as described in Section 7.2.1, the condition for catastrophic error propagation is that the generators have a *common polynomial factor* (of degree at least one). For example, Figure 7.20a illustrates a rate $\frac{1}{2}$, $K = 3$ encoder with upper polynomial $\mathbf{g}_1(X)$ and lower polynomial $\mathbf{g}_2(X)$, as follows:

$$\mathbf{g}_1(X) = 1 + X \tag{7.16}$$
$$\mathbf{g}_2(X) = 1 + X^2$$

The generators $\mathbf{g}_1(X)$ and $\mathbf{g}_2(X)$ have in common the polynomial factor $1 + X$, since

$$1 + X^2 = (1 + X)(1 + X)$$

Therefore, the encoder in Figure 7.20a can manifest *catastrophic error propagation.*

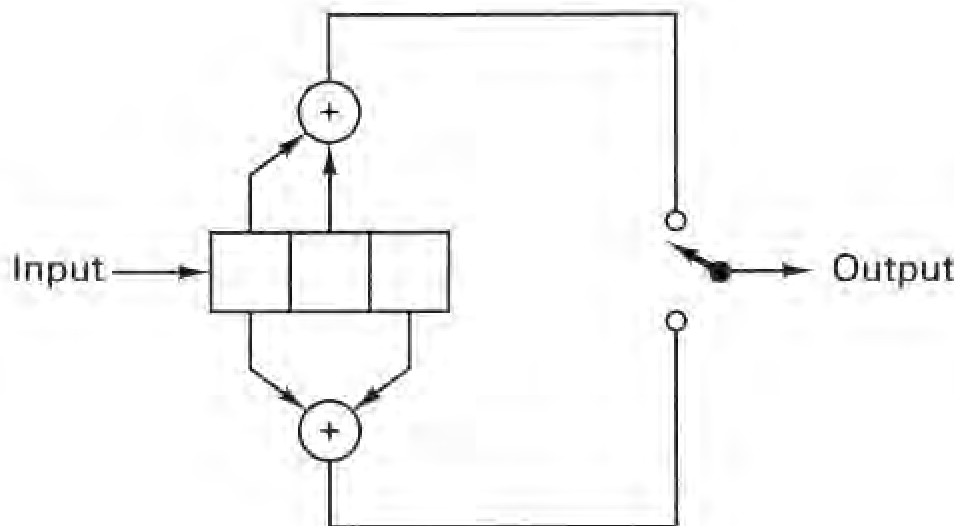

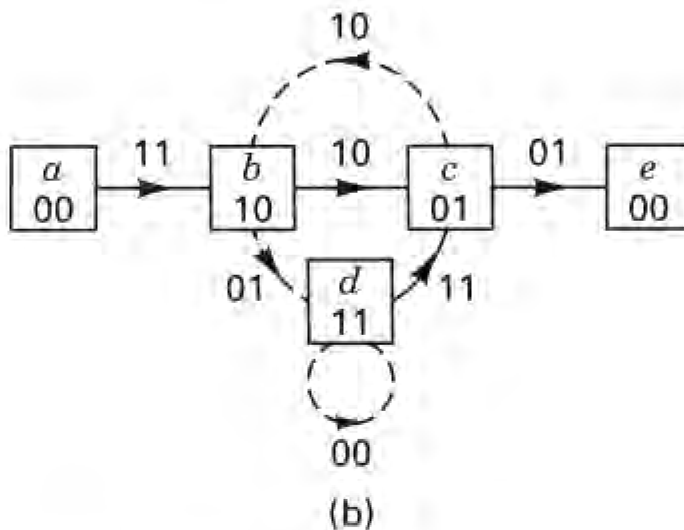

Figure 7.20 Encoder displaying catastrophic error propagation. (a) Encoder. (b) State diagram.

In terms of the state diagram for any-rate code, catastrophic errors can occur if, and only if, any closed-loop path in the diagram has zero weight (zero distance from the all-zeros path). To illustrate this, consider the example of Figure 7.20. The state diagram in Figure 7.20b is drawn with the state $a = 00$ node split into two nodes, a and e, as before. Assuming that the all-zeros path is the correct path, the incorrect path $a\ b\ d\ d \ldots d\ c\ e$ has exactly 6 ones, no matter how many times we go around the self-loop at node d. Thus for a BSC, for example, three channel errors may cause us to choose this incorrect path. An arbitrarily large number of errors (two plus the number of times the self-loop is traversed) can be made on such a path. We observe that for rate $1/n$ codes, if each adder in the encoder has an even number of connections, the self-loop corresponding to the all-ones data state will have zero weight, and consequently, *the code will be catastrophic.*

The only advantage of a systematic code, described earlier, is that it can never be catastrophic, since each closed loop must contain at least one branch generated by a nonzero input bit, and thus each closed loop must have a nonzero code symbol. However, it can be shown [19] that only a small fraction of nonsystematic codes (excluding those where all adders have an even number of taps) are catastrophic.

7.4.4 Performance Bounds for Convolutional Codes

The probability of bit error, P_B, for a binary convolutional code using hard-decision decoding can be shown [8] to be upper bounded as follows:

$$P_B \leq \left.\frac{dT(D,N)}{dN}\right|_{N=1,D=2\sqrt{p(1-p)}} \tag{7.17}$$

where p is the probability of channel symbol error. For the example of Figure 7.3, $T(D, N)$ is obtained from $T(D, L, N)$ by setting $L = 1$ in Equation (7.15).

$$T(D,N) = \frac{D^5 N}{1 - 2DN} \tag{7.18}$$

and

$$\left.\frac{dT(D,N)}{dN}\right|_{N=1} = \frac{D^5}{(1-2D)^2} \tag{7.19}$$

Combining Equations (7.17) and (7.19), we can write

$$P_B \leq \frac{\{2[p(1-p)]^{1/2}\}^5}{\{1 - 4[p(1-p)]^{1/2}\}^2} \tag{7.20}$$

For coherent BPSK modulation over an additive white Gaussian noise (AWGN) channel, it can be shown [8] that the bit error probability is bounded by

$$P_B \leq Q\left(\sqrt{2d_f \frac{E_c}{N_0}}\right) \exp\left(d_f \frac{E_c}{N_0}\right) \left.\frac{dT(D,N)}{dN}\right|_{N=1,D=\exp(-E_c/N_0)} \tag{7.21}$$

where

$E_c/N_0 = rE_b/N_0$

E_b/N_0 = ratio of information bit energy to noise power spectral density

E_c/N_0 = ratio of channel symbol energy to noise power spectral density

$r = k/n$ = rate of the code

and $Q(x)$ is defined in Equations (3.43) and (3.44) and tabulated in Table B.1. Therefore, for the rate $\frac{1}{2}$ code with free distance $d_f = 5$, in conjunction with coherent BPSK and hard-decision decoding, we can write

$$P_B \le Q\left(\sqrt{\frac{5E_b}{N_0}}\right)\exp\left(\frac{5E_b}{2N_0}\right)\frac{\exp(-5E_b/2N_0)}{[1-2\exp(-E_b/2N_0)]^2} \tag{7.22}$$

$$\le \frac{Q(\sqrt{5E_b/N_0})}{[1-2\exp(-E_b/2N_0)]^2}$$

7.4.5 Coding Gain

Coding gain, as presented in Equation (6.19), is defined as the reduction, usually expressed in decibels, in the required E_b/N_0 to achieve a specified error probability of the coded system over an uncoded system with the same modulation and channel characteristics. Table 7.2 lists an upper bound on the coding gains, compared to uncoded coherent BPSK, for several maximum free distance convolutional codes with constraint lengths varying from 3 to 9 over a Gaussian channel with hard-decision decoding. The table illustrates that it is possible to achieve significant coding gain even with a simple convolutional code. The actual coding gain will vary with the required bit error probability [20].

Table 7.3 lists the measured coding gains, compared to uncoded coherent BPSK, achieved with hardware implementation or computer simulation over a Gaussian channel with soft-decision decoding [21]. The uncoded E_b/N_0 is given in the leftmost column. From Table 7.3 we can see that coding gain increases as the bit error probability is decreased. However, the coding gain cannot increase indefinitely; it has an upper bound as shown in the table. This bound in decibels can be shown [21] to be

$$\text{coding gain} \le 10\log_{10}(rd_f) \tag{7.23}$$

where r is the code rate and d_f is the free distance. Examination of Table 7.3 also reveals that at $P_B = 10^{-7}$, for code rates of $\frac{1}{2}$ and $\frac{2}{3}$, the weaker codes tend to be closer to the upper bound than are the more powerful codes.

Typically, Viterbi decoding is used over binary input channels with either hard or 3-bit soft quantized outputs. The constraint lengths vary between 3 and 9, the code rate is rarely smaller than $\frac{1}{3}$, and the path memory is usually a few con-

TABLE 7.2 Coding Gain Upper Bounds for Some Convolutional Codes

Rate $\frac{1}{2}$ Codes			Rate $\frac{1}{3}$ Codes		
K	d_f	Upper Bound (dB)	K	d_f	Upper Bound (dB)
3	5	3.97	3	8	4.26
4	6	4.76	4	10	5.23
5	7	5.43	5	12	6.02
6	8	6.00	6	13	6.37
7	10	6.99	7	15	6.99
8	10	6.99	8	16	7.27
9	12	7.78	9	18	7.78

Source: V. K. Bhargava, D. Haccoun, R. Matyas, and P. Nuspl, *Digital Communications by Satellite,* John Wiley & Sons, Inc., New York, 1981.

straint lengths [12]. The path memory refers to the depth of the input bit history stored by the decoder. From the Viterbi decoding example in Section 7.3.4, one might question the notion of a fixed path memory. It seems from the example that the decoding of a branch word, at any arbitrary node, can take place as soon as there is only a single surviving branch at that node. That is true; however, to actually implement the decoder in this way would entail an extensive amount of processing to continually check when the branch word can be decoded. Instead, *a fixed delay is provided,* after which the branch word is decoded. It has been shown [12, 22] that tracing back from the state with the lowest state metric, over a fixed amount of path history (about 4 or 5 times the constraint length), is sufficient to limit the degradation from the optimum decoder performance to about 0.1 dB for the BSC and Gaussian channels. Typical error performance simulation results are shown in Figure 7.21 for Viterbi decoding with hard decision quantization [12]. Notice that each increment in constraint length improves the required E_b/N_0 by a factor of approximately 0.5 dB at $P_B = 10^{-5}$.

TABLE 7.3 Basic Coding Gain (dB) for Soft Decision Viterbi Decoding

Uncoded E_b/N_0 (dB)	Code Rate		$\frac{1}{3}$		$\frac{1}{2}$			$\frac{2}{3}$		$\frac{3}{4}$	
	P_B	K	7	8	5	6	7	6	8	6	9
6.8	10^{-3}		4.2	4.4	3.3	3.5	3.8	2.9	3.1	2.6	2.6
9.6	10^{-5}		5.7	5.9	4.3	4.6	5.1	4.2	4.6	3.6	4.2
11.3	10^{-7}		6.2	6.5	4.9	5.3	5.8	4.7	5.2	3.9	4.8
Upper bound			7.0	7.3	5.4	6.0	7.0	5.2	6.7	4.8	5.7

Source: I. M. Jacobs, "Practical Applications of Coding," *IEEE Trans. Inf. Theory,* vol. IT20, May 1974, pp. 305–310.

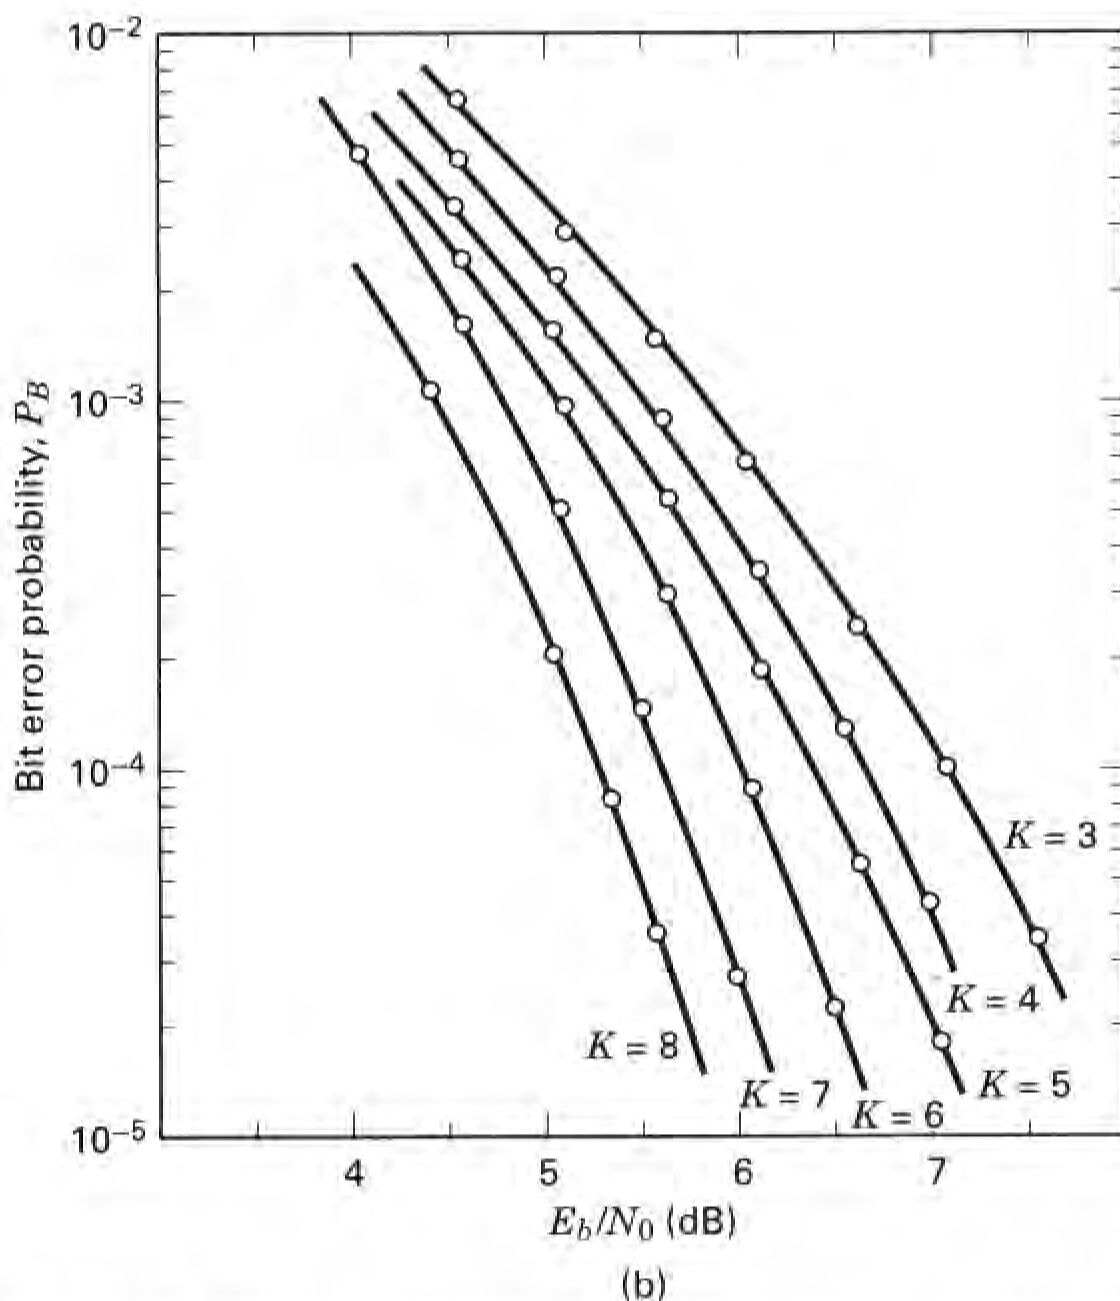

(b)

Figure 7.21 Bit error probability versus E_b/N_0 for rate $\frac{1}{2}$ codes using coherent BPSK over a BSC, Viterbi decoding, and a 32-bit path memory. (Reprinted with permission from J. A. Heller and I. M. Jacobs, "Viterbi Decoding for Satellite and Space Communication," *IEEE Trans. Commun. Technol.*, vol. COM19, no. 5, October 1971, Fig. 7, p. 84. © 1971 IEEE.)

7.4.6 Best Known Convolutional Codes

The connection vectors or polynomial generators of a convolutional code are usually selected based on the code's free distance properties. The first criterion is to select a code that does not have catastrophic error propagation and that has the maximum free distance for the given rate and constraint length. Then the number of paths at the free distance d_f, or the number of data bit errors the paths represent, should be minimized. The selection procedure can be further refined by considering the number of paths or bit errors at $d_f + 1$, at $d_f + 2$, and so on, until only one code or class of codes remains. A list of the best known codes of rate $\frac{1}{2}$, $K = 3$ to 9, and rate $\frac{1}{3}$, $K = 3$ to 8, based on this criterion was compiled by Odenwalder [3, 23] and is given in Table 7.4. The connection vectors in this table represent the pres-

ence or absence (1 or 0) of a tap connection on the corresponding stage of the convolutional encoder, the leftmost term corresponding to the leftmost stage of the encoder register. It is interesting to note that these connections can be inverted (leftmost and rightmost can be interchanged in the above description). Under the condition of Viterbi decoding, the inverted connections give rise to codes with identical distance properties, and hence identical performance, as those in Table 7.4.

TABLE 7.4 Optimum Short Constraint Length Convolutional Codes (Rate $\frac{1}{2}$ and Rate $\frac{1}{3}$)

Rate	Constraint Length	Free Distance	Code Vector
$\frac{1}{2}$	3	5	111 101
$\frac{1}{2}$	4	6	1111 1011
$\frac{1}{2}$	5	7	10111 11001
$\frac{1}{2}$	6	8	101111 110101
$\frac{1}{2}$	7	10	1001111 1101101
$\frac{1}{2}$	8	10	10011111 11100101
$\frac{1}{2}$	9	12	110101111 100011101
$\frac{1}{3}$	3	8	111 111 101
$\frac{1}{3}$	4	10	1111 1011 1101
$\frac{1}{3}$	5	12	11111 11011 10101
$\frac{1}{3}$	6	13	10111 110101 111001
$\frac{1}{3}$	7	15	1001111 1010111 1101101
$\frac{1}{3}$	8	16	11101111 10011011 10101001

Source: J. P. Odenwalder, *Error Control Coding Handbook*, Linkabit Corp., San Diego, Calif., July 15, 1976.

7.4.7 Convolutional Code Rate Trace-Off

7.4.7.1 Performance with Coherent PSK Signaling

The error-correcting capability of a coding scheme increases as the number of channel symbols n per information bit k increases, or the rate k/n decreases. However, the channel bandwidth and the decoder complexity both increase with n. The advantage of lower code rates when using convolutional codes with coherent PSK, is that the required E_b/N_0 is decreased (for a large range of code rates), permitting the transmission of higher data rates for a given amount of power, or permitting reduced power for a given data rate. Simulation studies have shown [16, 22] that for a fixed constraint length, a decrease in the code rate from $\frac{1}{2}$ to $\frac{1}{3}$ results in a reduction of the required E_b/N_0 of roughly 0.4 dB. However, the corresponding increase in decoder complexity is about 17%. For smaller values of code rate, the improvement in performance relative to the increased decoding complexity diminishes rapidly [22]. Eventually, a point is reached where further decrease in code rate is characterized by a reduction in coding gain. (See Section 9.7.7.2.)

7.4.7.2 Performance with Noncoherent Orthogonal Signaling

In contrast to PSK, there is an optimum code rate of about $\frac{1}{2}$ for noncoherent orthogonal signaling. Error performance at rates of $\frac{1}{3}$, $\frac{2}{3}$, and $\frac{3}{4}$ are each worse than those for rate $\frac{1}{2}$. For a fixed constraint length, the rate $\frac{1}{3}$, $\frac{2}{3}$, and $\frac{3}{4}$ codes typically degrade by about 0.25, 0.5, and 0.3 dB, respectively, relative to the rate $\frac{1}{2}$ performance [16].

7.4.8 Soft-Decision Viterbi Decoding

For a rate $\frac{1}{2}$ binary convolutional code system, the demodulator delivers two code symbols at a time to the decoder. For hard-decision (2-level) decoding, each pair of received code symbols can be depicted on a plane, as one of the corners of a square, as shown in Figure 7.22a. The corners are labeled with the binary numbers (0,0), (0,1), (1,0), and (1,1), representing the four possible hard-decision values that the two code symbols might have. For 8-level soft-decision decoding, each pair of code symbols can be similarly represented on an equally spaced 8-level by 8-level plane, as a point from the set of 64 points shown in Figure 7.22b. In this soft-decision case, the demodulator no longer delivers firm decisions; it delivers quantized noisy signals (soft decisions).

The primary difference between hard-decision and soft-decision Viterbi decoding, is that the soft-decision algorithm cannot use a Hamming distance metric because of its limited resolution. A distance metric with the needed resolution is Euclidean distance, and to facilitate its use, the binary numbers of 1 and 0 are transformed to the octal numbers 7 and 0, respectively. This can be seen in Figure 7.22c, where the corners of the square have been re-labeled accordingly; this allows us to use a pair of integers, each in the range of 0 to 7, for describing any point in the 64-point set. Also shown in Figure 7.22c is the point 5,4, representing an example of a pair of noisy code-symbol values that might stem from a

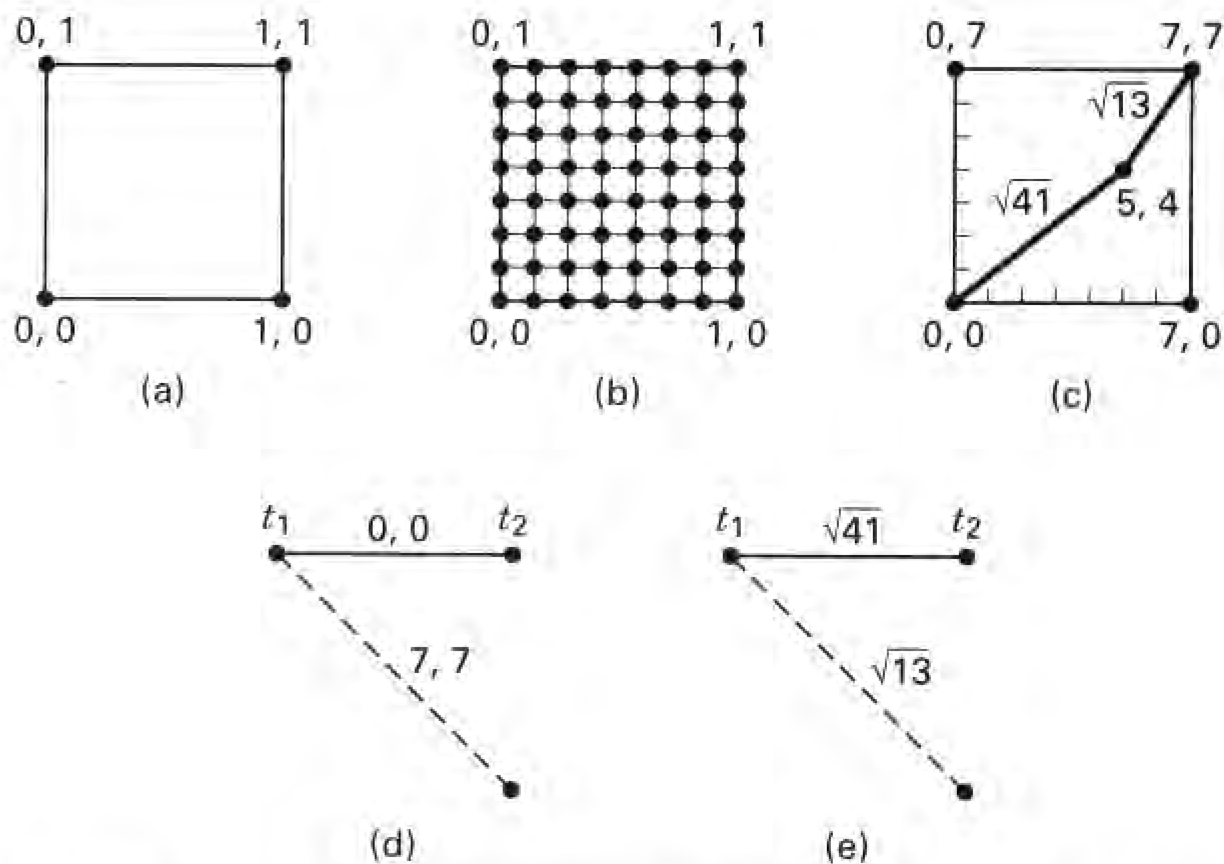

Figure 7.22 (a) Hard-decision plane (b) 8-level by 8-level soft-decision plane (c) Example of soft code symbols (d) Encoding trellis section (e) Decoding trellis section.

demodulator. Imagine that the square in Figure 7.22c has coordinates x and y. Then, what is the Euclidean distance between the noisy point 5,4 and the noiseless point 0,0? It is $\sqrt{(5-0)^2+(4-0)^2}=\sqrt{41}$. Similarly, if we ask what is the Euclidean distance between the noisy point 5,4 and the noiseless point 7,7? It is $\sqrt{(5-7)^2+(4-7)^2}=\sqrt{13}$.

Soft-decision Viterbi decoding, for the most part, proceeds in the same way as hard-decision decoding (as described in Sections 7.3.4 and 7.3.5). The only difference is that Hamming distances are not used. Consider how soft-decision decoding is performed with the use of Euclidean distances. Figure 7.22d shows the first section of an encoding trellis, originally presented in Figure 7.7, with the branch words transformed from binary to octal. Suppose that a pair of soft-decision code symbols with values 5,4 arrives at a decoder during the first transition interval. Figure 7.22e shows the first section of a decoding trellis. The metric ($\sqrt{41}$), representing the Euclidean distance between the arriving 5,4 and the 0,0 branch word, is placed on the solid line. Similarly, the metric ($\sqrt{13}$), representing the Euclidean distance between the arriving 5,4 and the 7,7 code symbols, is placed on the dashed line. The rest of the task, pruning the trellis in search of a common stem, proceeds in the same way as hard-decision decoding. Note that in a real convolutional decoding chip, the Euclidean distance is not actually used for a soft-decision metric; instead, a monotonic metric that has similar properties and is easier to implement is used. An example of such a metric is the Euclidean distance-squared, in which case the square-root operation shown above is eliminated. Further, if the binary code symbols are represented with bipolar

values, then the inner-product metric in Equation (7.9) can be used. With such a metric, we would seek maximum correlation rather than minimum distance.

7.5 OTHER CONVOLUTIONAL DECODING ALGORITHMS

7.5.1 Sequential Decoding

Prior to the discovery of an optimum algorithm by Viterbi, other algorithms had been proposed for decoding convolutional codes. The earliest was the *sequential decoding algorithm,* originally proposed by Wozencraft [24, 25] and subsequently modified by Fano [2]. A sequential decoder works by generating hypotheses about the transmitted codeword sequence; it computes a metric between these hypotheses and the received signal. It goes forward as long as the metric indicates that its choices are likely; otherwise, it goes backward, changing hypotheses until, through a systematic trial-and-error search, it finds a likely hypothesis. Sequential decoders can be implemented to work with hard or soft decisions, but soft decisions are usually avoided because they greatly increase the amount of the required storage and the complexity of the computations.

Consider that using the encoder shown in Figure 7.3, a sequence **m** = 1 1 0 1 1 is encoded into the codeword sequence **U** = 1 1 0 1 0 1 0 0 0 1, as shown in Example 7.1. Assume that the received sequence **Z** is, in fact, a *correct* rendition of **U**. The decoder has available a replica of the encoder code tree, shown in Figure 7.6, and can use the received sequence **Z** to penetrate the tree. The decoder starts at the time t_1 node of the tree and generates both paths leaving that node. The decoder follows that path which agrees with the received n code symbols. At the next level in the tree, the decoder again generates both paths leaving that node, and follows the path agreeing with the second group of n code symbols. Proceeding in this manner, the decoder quickly penetrates the tree.

Suppose, however, that the received sequence **Z** is a *corrupted* version of **U**. The decoder starts at the time t_1 node of the code tree and generates both paths leading from that node. If the received n code symbols coincide with one of the generated paths, the decoder follows that path. If there is not agreement, the decoder follows the *most likely path* but keeps a cumulative count on the number of disagreements between the received symbols and the branch words on the path being followed. If two branches appear equally likely, the receiver uses an arbitrary rule, such as following the zero input path. At each new level in the tree, the decoder generates new branches and compares them with the next set of n received code symbols. The search continues to penetrate the tree along the most likely path and maintains the cumulative disagreement count.

If the disagreement count exceeds a certain number (which may increase as we penetrate the tree), the decoder decides that it is on an incorrect path, backs out of the path, and tries another. The decoder keeps track of the discarded pathways to avoid repeating any path excursions. For example, assume that the encoder in Figure 7.3 is used to encode the message sequence **m** = 1 1 0 1 1 into the codeword sequence **U** as shown in Example 7.1. Suppose that the fourth and seventh bits of the transmitted sequence **U** are received in error, such that:

Time:		t_1	t_2	t_3	t_4	t_5
Message sequence:	**m** =	1	1	0	1	1
Transmitted sequence:	**U** =	1 1	0 1	0 1	0 0	0 1
Received sequence:	**Z** =	1 1	0 0	0 1	1 0	0 1

Let us follow the decoder path trajectory with the aid of Figure 7.23. Assume that a cumulative path disagreement count of 3 is the criterion for backing up and trying an alternative path. On Figure 7.23 the numbers along the path trajectory represent the current disagreement count.

1. At time t_1 we receive symbols 11 and compare them with the branch words leaving the first node.
2. The most likely branch is the one with branch word 11 (corresponding to an input bit one or downward branching), so the decoder decides that input bit one is the correct decoding, and moves to the next level.
3. At time t_2, the decoder receives symbols 00 and compares them with the available branch words 10 and 01 at this second level.
4. There is no "best" path, so the decoder arbitrarily takes the input bit zero (or branch word 10) path, and the disagreement count registers a disagreement of 1.
5. At time t_3, the decoder receives symbols 01 and compares them with the available branch words 11 and 00 at this third level.
6. Again, there is no best path, so the decoder arbitrarily takes the input zero (or branch word 11) path, and the disagreement count is increased to 2.
7. At time t_4, the decoder receives symbols 10 and compares them with the available branch words 00 and 11 at this fourth level.
8. Again, there is no best path, so the decoder takes the input bit zero (or branch word 00) path, and the disagreement count is increased to 3.
9. But a disagreement count of 3 is the turnaround criterion, so the decoder "backs out" and tries the alternative path. The disagreement counter is reset to 2.
10. The alternative path is the input bit one (or branch word 11) path at the t_4 level. The decoder tries this, but compared to the received symbols 10, there is still a disagreement of 1, and the counter is reset to 3.
11. But, 3 being the turnaround criterion, the decoder backs out of this path, and the counter is reset to 2. All of the alternatives have now been traversed at this t_4 level, so the decoder returns to the node at t_3, and resets the counter to 1.
12. At the t_3 node, the decoder compares the symbols received at time t_3, namely 01, with the untried 00 path. There is a disagreement of 1, and the counter is increased to 2.

t_1 t_2 t_3 t_4 t_5 t_6

$\mathbf{Z}$ = 11 00 01 10 01

Figure 7.23 Sequential decoding example.

13. At the t_4 node, the decoder follows the branch word 10 that matches its t_4 code symbols of 10. The counter remains unchanged at 2.
14. At the t_5 node, there is no best path, so the decoder follows the upper branch, as is the rule, and the counter is increased to 3.
15. At this count, the decoder backs up, resets the counter to 2, and tries the alternative path at node t_5. Since the alternate branch word is 00, there is a disagreement of 1 with the received code symbols 01 at time t_5, and the counter is again increased to 3.
16. The decoder backs out of this path, and the counter is reset to 2. All of the alternatives have now been traversed at this t_5 level, so the decoder returns to the node at t_4 and resets the counter to 1.
17. The decoder tries the alternative path at t_4, which raises the metric to 3 since there is a disagreement in two positions of the branch word. This time the decoder must back up all the way to the time t_2 node because all of the other paths at higher levels have been tried. The counter is now decremented to zero.
18. At the t_2 node, the decoder now follows the branch word 01, and because there is a disagreement of 1 with the received code symbols 00 at time t_2, the counter is increased to 1.

The decoder continues in this way. As shown in Figure 7.23, the final path, which has not increased the counter to its turnaround criterion, yields the correctly decoded message sequence, 1 1 0 1 1. Sequential decoding can be viewed as a trial-and-error technique for searching out the correct path in the code tree. It performs the search in a sequential manner, always operating on just a single path at a time. If an incorrect decision is made, subsequent extensions of the path will be wrong. The decoder can eventually recognize its error by monitoring the path metric. The algorithm is similar to the case of an automobile traveler following a road map. As long as the traveler recognizes that the passing landmarks correspond to those on the map, he continues on the path. When he notices strange landmarks (an increase in his dissimilarity metric) the traveler eventually assumes that he is on an incorrect road, and he backs up to a point where he can now recognize the landmarks (his metric returns to an acceptable range). He then tries an alternative road.

7.5.2 Comparisons and Limitations of Viterbi and Sequential Decoding

The major drawback of the Viterbi algorithm is that while error probability decreases exponentially with constraint length, the number of code states, and consequently decoder complexity, *grows exponentially with constraint length.* On the other hand, the computational complexity of the Viterbi algorithm is independent of channel characteristics (compared to hard-decision decoding, soft-decision decoding requires only a trivial increase in the number of computations). Sequential decoding achieves asymptotically the same error probability as maximum likelihood decoding but without searching all possible states. In fact, with sequential de-

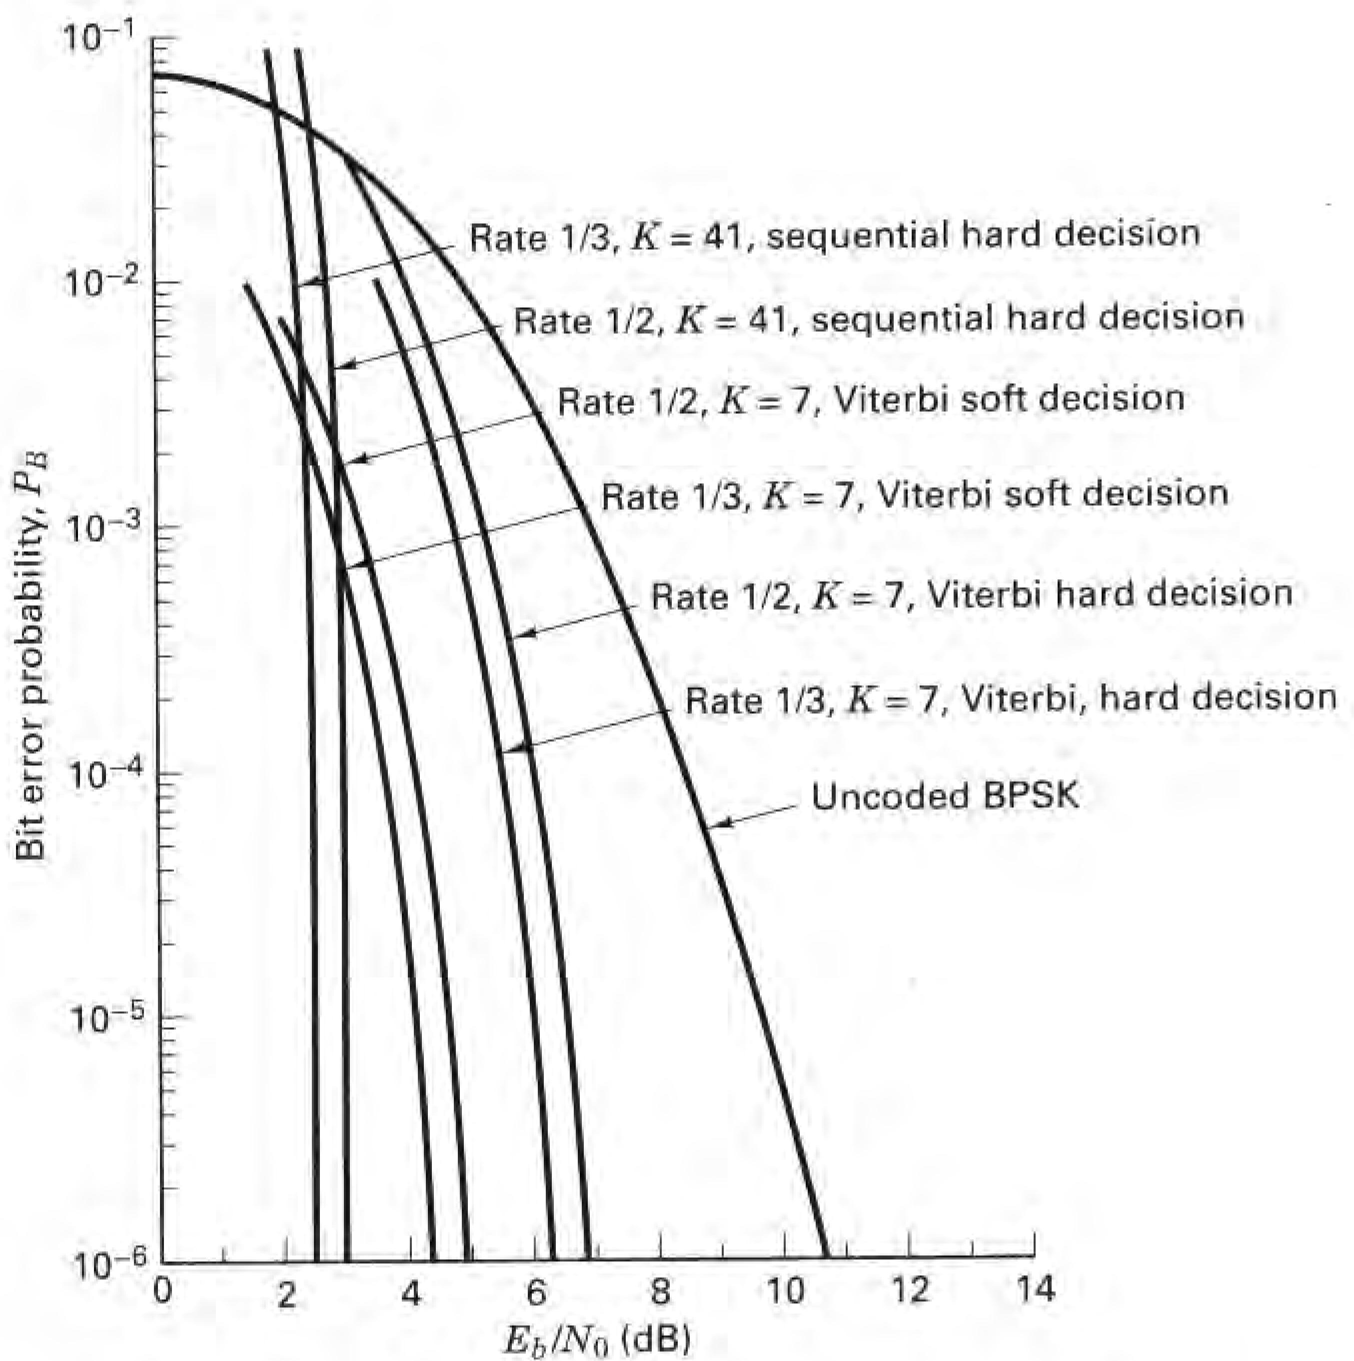

Figure 7.24 Bit error performance for various Viterbi and sequential decoding schemes using coherent BPSK over an AWGN channel. (Reprinted with permission from J. K. Omura and B. K. Levitt, "Coded Error Probability Evaluation for Antijam Communication Systems," *IEEE Trans. Commun.*, vol. COM30, no. 5, May 1982, Fig. 4, p. 900. © 1982 IEEE.)

coding the number of states searched is essentially *independent of constraint length*, thus making it possible to use very large ($K = 41$) constraint lengths. This is an important factor in providing such low error probabilities. The major drawback of sequential decoding is that the number of state metrics searched is a random variable. For sequential decoding, the expected number of poor hypotheses and backward searches is a function of the channel SNR. With a low SNR, more hypotheses must be tried than with a high SNR. Because of this variability in computational load, buffers must be provided to store the arriving sequences. Under low SNR, the received sequences must be buffered while the decoder is laboring to find a likely hypothesis. If the average symbol arrival rate exceeds the average symbol decode rate, the buffer will overflow, no matter how large it is, causing a loss of

data. The sequential decoder typically puts out error-free data until the buffer overflows, at which time the decoder has to go through a recovery procedure. The buffer overflow threshold is a very sensitive function of SNR. Therefore, an important part of a sequential decoder specification is the *probability of buffer overflow.*

In Figure 7.24, some typical P_B versus E_b/N_0 curves for these two popular solutions to the convolutional decoding problem, Viterbi decoding and sequential decoding, illustrate their comparative performance using coherent BPSK over an AWGN channel. The curves compare Viterbi decoding (rates $\frac{1}{2}$ and $\frac{1}{3}$ hard decision, $K = 7$) versus Viterbi decoding (rates $\frac{1}{2}$ and $\frac{1}{3}$ soft decision, $K = 7$) versus sequential decoding (rates $\frac{1}{2}$ and $\frac{1}{3}$ hard decision, $K = 41$). One can see from Figure 7.24 that coding gains of approximately 8 dB at $P_B = 10^{-6}$ can be achieved with sequential decoders. Since the work of Shannon [26] foretold the potential of approximately 11 dB of coding gain compared to uncoded BPSK, it appears that the major portion of what is theoretically possible can already be accomplished.

7.5.3 Feedback Decoding

A *feedback decoder* makes a hard decision on the data bit at stage j based on metrics computed from stages $j, j + 1, \ldots, j + m$, where m is a preselected positive integer. *Look-ahead length, L,* is defined as $L = m + 1$, the number of received code symbols, expressed in terms of the corresponding number of encoder input bits that are used to decode an information bit. The decision of whether the data bit is zero or one depends on which branch the minimum Hamming distance path traverses in the *look-ahead window* from stage j to stage $j + m$. The detailed operation is best understood in terms of a specific example. Let us consider the use of a feedback decoder for the rate $\frac{1}{2}$ convolutional code shown in Figure 7.3. Figure 7.25 illustrates the tree diagram and the operation of the feedback decoder for $L = 3$. That is, in decoding the bit at branch j, the decoder considers the paths at branches j, $j + 1$, and $j + 2$.

Beginning with the first branch, the decoder computes 2^L or eight cumulative Hamming path metrics and decides that the bit for the first branch is zero if the minimum distance path is contained in the upper part of the tree, and decides one if the minimum distance path is in the lower part of the tree. Assume that the received sequence is $\mathbf{Z} = 1\ 1\ 0\ 0\ 0\ 1\ 0\ 0\ 0\ 1$. We now examine the eight paths from time t_1 through time t_3 in the block marked A in Figure 7.24, and compute metrics comparing these eight paths with the first six received code symbols (three branches deep times two symbols per branch). Listing the Hamming cumulative path metrics (starting from the top path), we see that they are

Upper-half metrics: 3, 3, 6, 4

Lower-half metrics: 2, 2, 1, 3

We see that the minimum metric is contained in the lower part of the tree. Therefore, the first decoded bit is one (characterized by a downward movement on the tree). The next step is to extend the lower part of the tree (the part that survived)

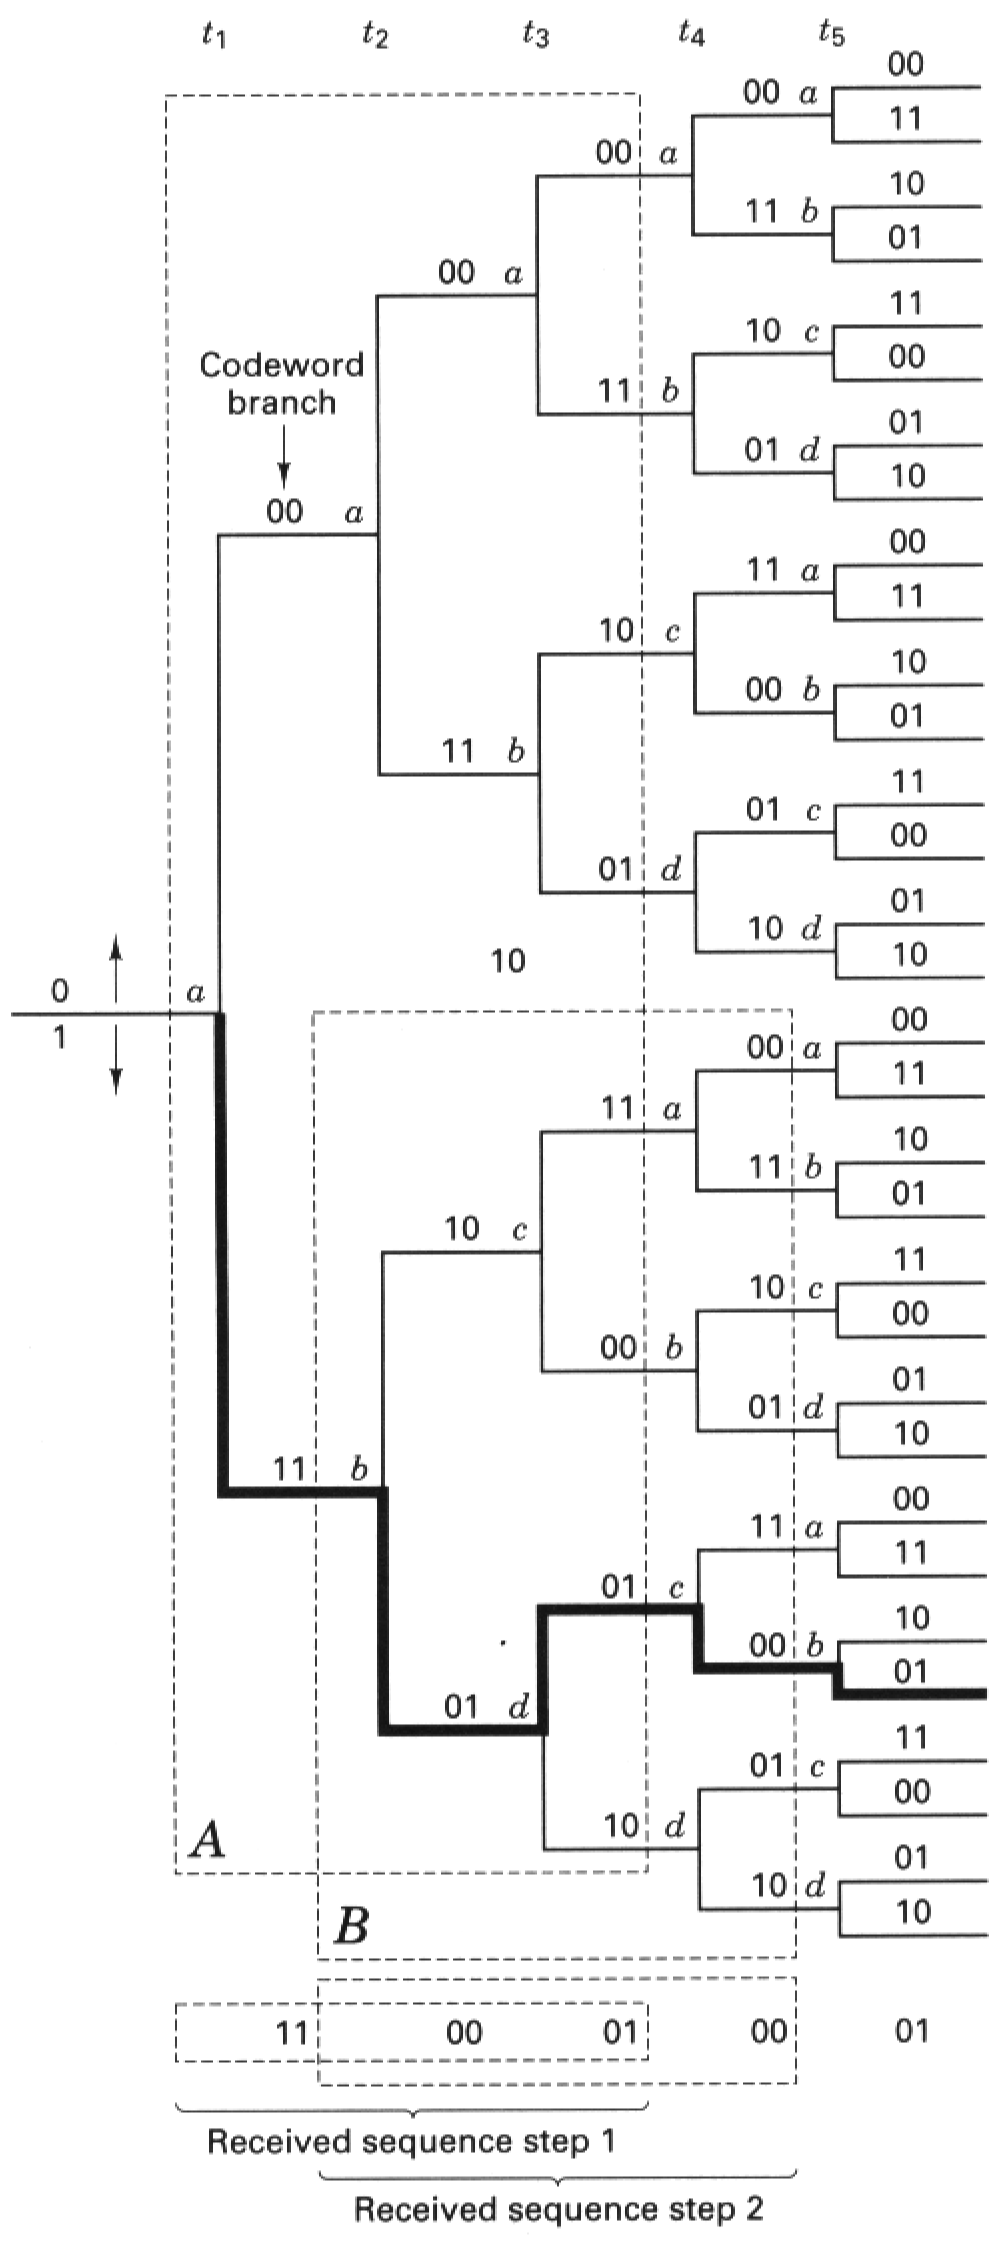

Figure 7.25 Feedback decoding example.

one stage deeper, and again compute eight metrics, this time from t_2 through t_4. Having decoded the first two code symbols, we now slide over two code symbols to the right and again compute the path metrics for six code symbols. This takes place in the block marked B in Figure 7.25. Again, listing the metrics from top path to bottom path, we find that they are

Upper-half metrics: 2, 4, 3, 3

Lower-half metrics: 3, 1, 4, 4

For the assumed received sequence, the minimum metric is found in the lower half of block B. Therefore, the second decoded bit is one.

The same procedure continues until the entire message is decoded. The decoder is called a *feedback decoder* because the detection decisions are *fed back* to the decoder in determining the subset of code paths that are to be considered next. On the BSC, the feedback decoder can perform nearly as well as the Viterbi decoder [17] in that it can correct all the more probable error patterns, namely all those of weight $(d_f - 1)/2$ or less, where d_f is the free distance of the code. An important design parameter for feedback convolutional decoders is L, the look-ahead length. Increasing L increases the coding gain but also increases the decoder implementation complexity.

7.6 CONCLUSION

In the last decade, coding emphasis has been in the area of convolutional codes since in almost every application, convolutional codes outperform block codes for the same implementation complexity of the encoder–decoder. For satellite communication channels, forward error correction techniques can easily reduce the required SNR for a specified error performance by 5 to 6 dB. This coding gain can translate directly into an equivalent reduction in required satellite effective radiated power (EIRP), with consequently reduced satellite weight and cost.

In this chapter we have outlined the essential structural difference between block codes and convolutional codes—the fact that rate $1/n$ convolutional codes have a memory of the prior $K - 1$ bits, where K is the encoder constraint length. With such memory, the encoding of each input data bit not only depends on the value of that bit but on the values of the $K - 1$ input bits that precede it. We presented the decoding problem in the context of the maximum likelihood algorithm, examining all the candidate codeword sequences which could possibly be created by the encoder, and selecting the one that appears statistically most likely; the decision is based on a distance metric for the received code symbols. The error performance analysis of convolutional codes is more complicated than the simple binomial expansion describing the error performance of many block codes. We laid out the concept of free distance, and we presented the relationship between free distance and error performance in terms of bounds. We also described the basic

idea behind sequential decoding and feedback decoding and showed some comparative performance curves and tables for various coding schemes.

REFERENCES

1. Gallager, R. G., *Information Theory and Reliable Communication,* John Wiley & Sons, Inc., New York, 1968.
2. Fano, R. M., "A Heuristic Discussion of Probabilistic Decoding," IRE *Trans. Inf. Theory,* vol. IT9, no. 2, 1963, pp. 64–74.
3. Odenwalder, J. P., *Optimal Decoding of Convolutional Codes,* Ph.D. dissertation, University of California, Los Angeles, 1970.
4. Curry, S. J., *Selection of Convolutional Codes Having Large Free Distance,* Ph.D. dissertation, University of California, Los Angeles, 1971.
5. Larsen, K. J., "Short Convolutional Codes with Maximal Free Distance for Rates $\frac{1}{2}$, $\frac{1}{3}$, and $\frac{1}{4}$," *IEEE Trans. Inf. Theory,* vol. IT19, no. 3, 1973, pp. 371–372.
6. Lin, S., and Costello, D. J., Jr., *Error Control Coding: Fundamentals and Applications,* Prentice-Hall, Inc., Englewood Cliffs, N.J., 1983.
7. Forney, G. D., Jr., "Convolutional Codes: I. Algebraic Structure," *IEEE Trans. Inf. Theory,* vol. IT16, no. 6, Nov. 1970, pp. 720–738.
8. Viterbi, A., "Convolutional Codes and Their Performance in Communication Systems," *IEEE Trans. Commun. Technol.,* vol. COM 19, no. 5, Oct. 1971, pp. 751–772.
9. Forney, G. D., Jr., and Bower, E. K., "A High Speed Sequential Decoder: Prototype Design and Test," *IEEE Trans. Commun. Technol.,* vol. COM19, no. 5, Oct. 1971, pp. 821–835.
10. Jelinek, F., "Fast Sequential Decoding Algorithm Using a Stack," *IBM J. Res. Dev.,* vol. 13, Nov. 1969, pp. 675–685.
11. Massey, J. L., *Threshold Decoding,* The MIT Press, Cambridge, Mass., 1963.
12. Heller, J. A., and Jacobs, I. W., "Viterbi Decoding for Satellite and Space Communication," *IEEE Trans. Commun. Technol.,* vol. COM19, no. 5, October 1971, pp. 835–848.
13. Viterbi, A. J., "Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm," *IEEE Trans. Inf. Theory,* vol. IT13, April 1967, pp. 260–269.
14. Omura, J. K., "On the Viterbi Decoding Algorithm" (correspondence), *IEEE Trans. Inf. Theory,* vol. IT15, Jan. 1969, pp. 177–179.
15. Mason, S. J., and Zimmerman, H. J., *Electronic Circuits, Signals, and Systems,* John Wiley & Sons, Inc., New York, 1960.
16. Clark, G. C., Jr., and Cain, J. B., *Error-Correction Coding for Digital Communications,* Plenum Press, New York, 1981.
17. Viterbi, A. J., and Omura, J. K., *Principles of Digital Communication and Coding,* McGraw-Hill Book Company, New York, 1979.
18. Massey, J. L., and Sain, M. K., "Inverse of Linear Sequential Circuits," *IEEE Trans. Comput.,* vol. C17, Apr. 1968, pp. 330–337.
19. Rosenberg, W. J., *Structural Properties of Convolutional Codes,* Ph.D. dissertation, University of California, Los Angeles, 1971.

20. Bhargava, V. K., Haccoun, D., Matyas, R., and Nuspl, P., *Digital Communications by Satellite,* John Wiley & Sons, Inc., New York, 1981.

21. Jacobs, I. M., "Practical Applications of Coding," *IEEE Trans. Inf. Theory,* vol. IT20, May 1974, pp. 305–310.

22. Linkabit Corporation, "Coding Systems Study for High Data Rate Telemetry Links," *NASA Ames Res. Center, Final Rep. CR-114278,* Contract NAS-2-6-24, Moffett Field, Calif., 1970.

23. Odenwalder, J. P., *Error Control Coding Handbook,* Linkabit Corporation, San Diego, Calif., July 15, 1976.

24. Wozencraft, J. M., "Sequential Decoding for Reliable Communication," *IRE Natl. Conv. Rec.,* vol. 5, pt. 2, 1957, pp. 11–25.

25. Wozencraft, J. M., and Reiffen, B., *Sequential Decoding,* The MIT Press, Cambridge, Mass., 1961.

26. Shannon, C. E., "A Mathematical Theory of Communication," *Bell Syst. Tech. J.,* vol. 27, 1948, pp. 379–423, 623–656.

PROBLEMS

7.1. Draw the state diagram, tree diagram, and trellis diagram for the $K = 3$, rate $\frac{1}{3}$ code generated by

$$\mathbf{g}_1(X) = X + X^2$$

$$\mathbf{g}_2(X) = 1 + X$$

$$\mathbf{g}_3(X) = 1 + X + X^2$$

7.2. Given a $K = 3$, rate $\frac{1}{2}$, binary convolutional code with the partially completed state diagram shown in Figure P7.1, find the complete state diagram and sketch a diagram for the encoder.

7.3. Draw the state diagram, tree diagram, and trellis diagram for the convolutional encoder characterized by the block diagram in Figure P7.2.

7.4. Suppose that you were trying to find the quickest way to get from London to Vienna by boat or train. The diagram in Figure P7.3 was constructed from various schedules. The labels on each path are travel times. Using the Viterbi algorithm, find the fastest

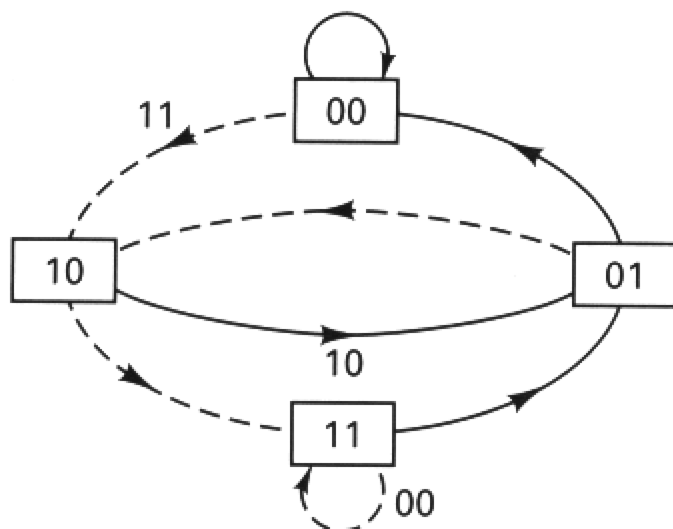

Figure P7.1

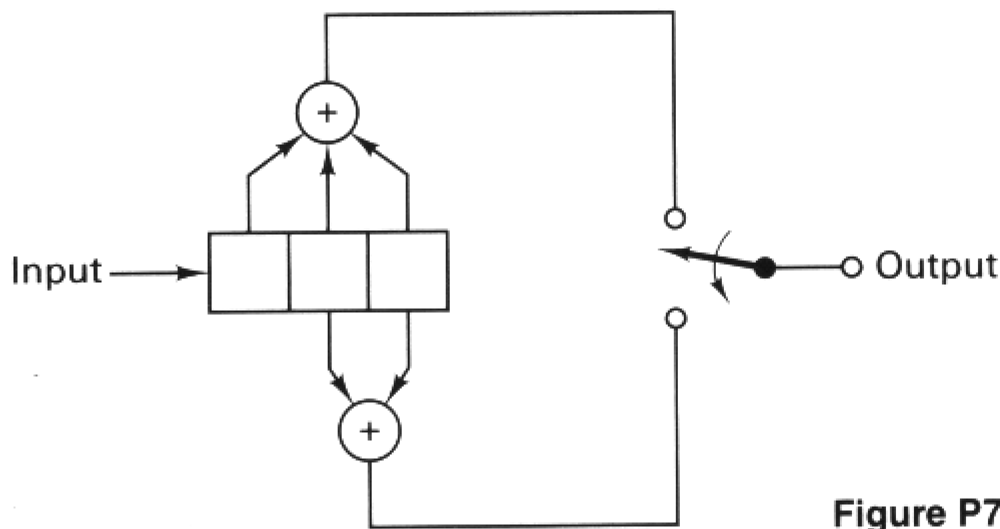

Figure P7.2

route from London to Vienna. In a general sense, explain how the algorithm works, what calculations must be made, and what information must be retained in the memory used by the algorithm.

7.5. Consider the convolutional encoder shown in Figure P7.4.
  (a) Write the connection vectors and polynomials for this encoder.
  (b) Draw the state diagram, tree diagram, and trellis diagram.

7.6. What is the impulse response of the encoder of Problem 7.5? Using the impulse response, determine the output sequence when the input is 1 0 1. Verify by using the generator polynomials.

7.7. Does the encoder of Problem 7.5 exhibit the properties of catastrophic error propagation? Justify your answer with an example.

7.8. Find the free distance of the encoder of Problem 7.3 by the transfer function method.

7.9. Let the codewords of a coding scheme be

$$a = 0\,0\,0\,0\,0\,0$$
$$b = 1\,0\,1\,0\,1\,0$$
$$c = 0\,1\,0\,1\,0\,1$$
$$d = 1\,1\,1\,1\,1\,1$$

If the received sequence over a binary symmetric channel is 1 1 1 0 1 0 and a maximum likelihood decoder is used, what will be the decoded symbol?

7.10. Consider that the $K = 3$, rate $\frac{1}{2}$ encoder of Figure 7.3 is used over a binary symmetric channel (BSC). Assume that the initial encoder state is the 00 state. At the output of the BSC, the sequence $\mathbf{Z} = (1\,1\,0\,0\,0\,0\,1\,0\,1\,1$ rest all "0") is received.

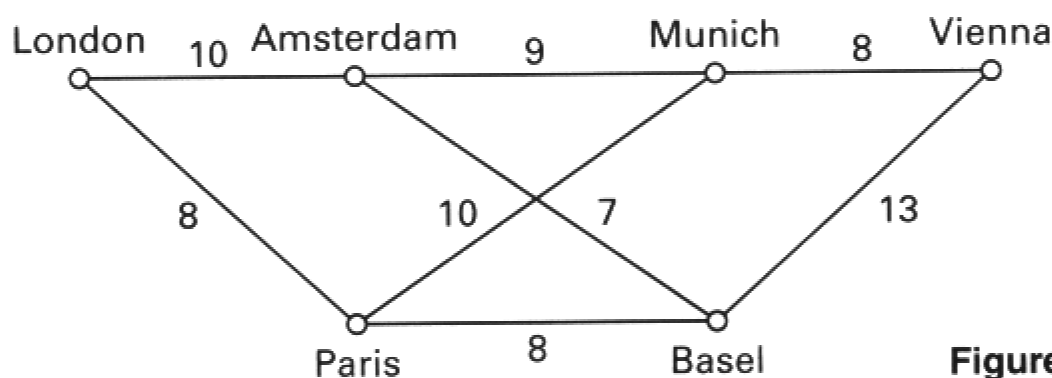

Figure P7.3

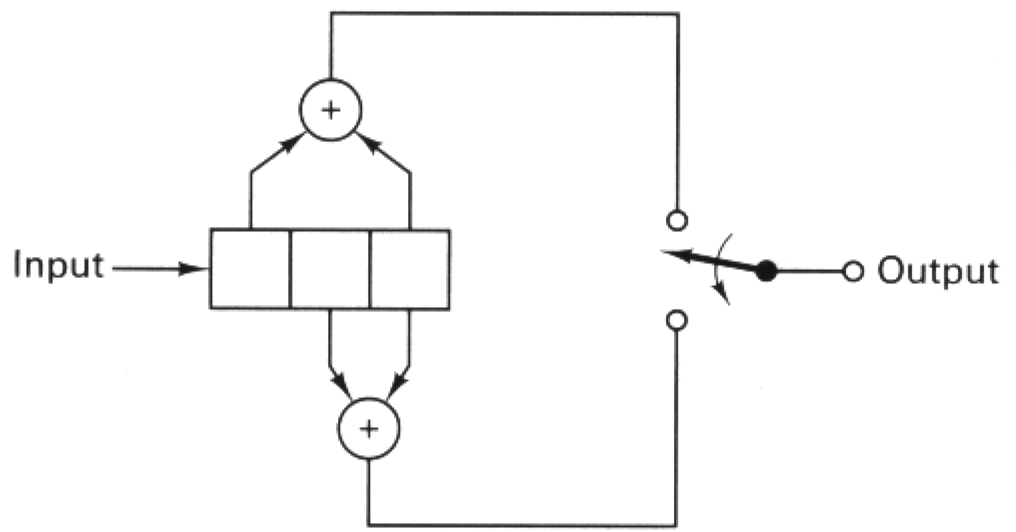

Figure P7.4

(a) Find the maximum likelihood path through the trellis diagram, and determine the first 5 decoded information bits. If a tie occurs between any two merged paths, choose the upper branch entering the particular state.

(b) Identify any channel bits in **Z** that were inverted by the channel during transmission.

7.11. Determine which of the following rate $\frac{1}{2}$ codes are catastrophic.

(a) $\mathbf{g}_1(X) = X^2, \quad \mathbf{g}_2(X) = 1 + X + X^3$

(b) $\mathbf{g}_1(X) = 1 + X^2, \quad \mathbf{g}_2(X) = 1 + X^3$

(c) $\mathbf{g}_1(X) = 1 + X + X^2, \quad \mathbf{g}_2(X) = 1 + X + X^3 + X^4$

(d) $\mathbf{g}_1(X) = 1 + X + X^3 + X^4, \quad \mathbf{g}_2(X) = 1 + X^2 + X^4$

(e) $\mathbf{g}_1(X) = 1 + X^4 + X^6 + X^7, \quad \mathbf{g}_2(X) = 1 + X^3 + X^4$

(f) $\mathbf{g}_1(X) = 1 + X^3 + X^4, \quad \mathbf{g}_2(X) = 1 + X + X^2 + X^4$

7.12. (a) Consider a coherently detected BPSK signal encoded with the encoder shown in Figure 7.3. Find an upper bound on the bit error probability, P_B, if the available E_b/N_0 is 6 dB. Assume hard decision decoding.

(b) Compare P_B with the uncoded case and calculate the improvement factor.

7.13. Using sequential decoding, illustrate the path along the tree diagram shown in Figure 7.22 when the received sequence is 0 1 1 1 0 0 0 1 1 1. The backup criterion is three disagreements.

7.14. Repeat the decoding example of Problem 7.13 using feedback decoding, with a look-ahead length of 3. In the event of a tie, select the upper half of the tree.

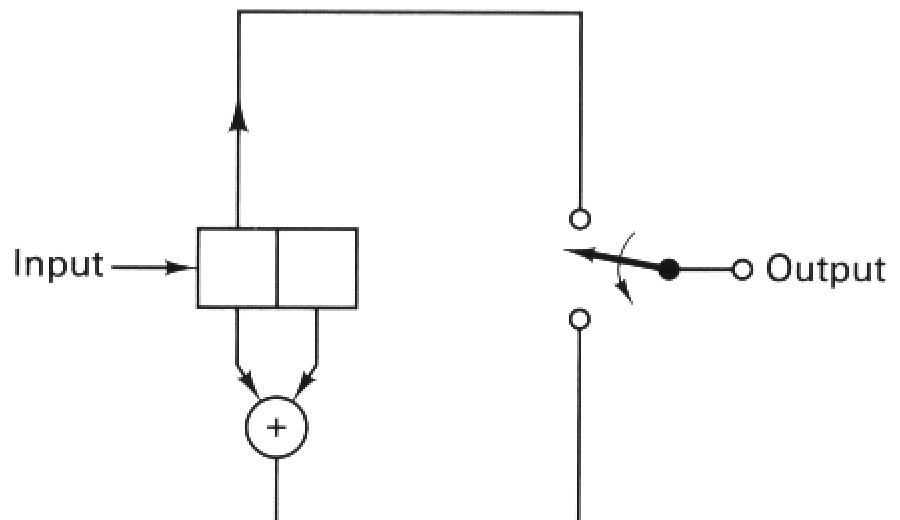

Figure P7.5

7.15. Figure P7.5 depicts a constraint length 2 convolutional encoder.

(a) Draw the state diagram, tree diagram, and trellis diagram.

(b) Assume that a received message from this encoder is 1 1 0 0 1 0. Use a feedback decoding algorithm with a look-ahead length of 2 to decode the coded message sequence.

7.16. Using the branch word information on the encoder trellis of Figure 7.7, decode the sequence $\mathbf{Z}$ = (01 11 00 01 11 rest all "0"), using hard-decision Viterbi decoding.

7.17. Consider the rate $\frac{2}{3}$ convolutional encoder shown in Figure P7.6. In this encoder, $k = 2$ bits at a time are shifted into the encoder and $n = 3$ bits are generated at the encoder output. There are $kK = 4$ stages in the register, and the constraint length is $K = 2$ in units of 2-bit bytes. The state of the encoder is defined as the contents of the rightmost $K - 1$ k-tuple stages. Draw the state diagram, the tree diagram, and the trellis diagram.

7.18. Find the ratio of the predetection signal-to-noise spectral density, P_r/N_0, in decibels, required to yield a decoded data rate of 1 Mbit/s with error probability of 10^{-5}. Assume binary noncoherent FSK modulation. Also, assume convolutional encoding with the decoder relationship

$$P_B = 2000\, p_c^4$$

where p_c and P_B are bit error probabilities into and out of the decoder, respectively.

7.19. Using Table 7.4, devise a $K = 4$, rate $\frac{1}{2}$ binary convolutional encoder.

(a) Draw the circuit.

(b) Draw the encoding trellis showing its states and branch words.

(c) Configure the cells that would be implemented in an ACS algorithm.

7.20. For the $K = 3$, rate $\frac{1}{2}$ code described by the encoder circuit of Figure 7.3, perform soft-decision decoding for the following demodulated sequence. The signals are 8-level quantized integers in the range of 0 to 7. The level 0 represents the perfect binary 0, and the level 7 represents the perfect binary 1. If the digits into the decoder are: 6, 7, 5, 3, 1, 0, 1, 1, 2, where the leftmost digit is the earliest, use a decoding trellis

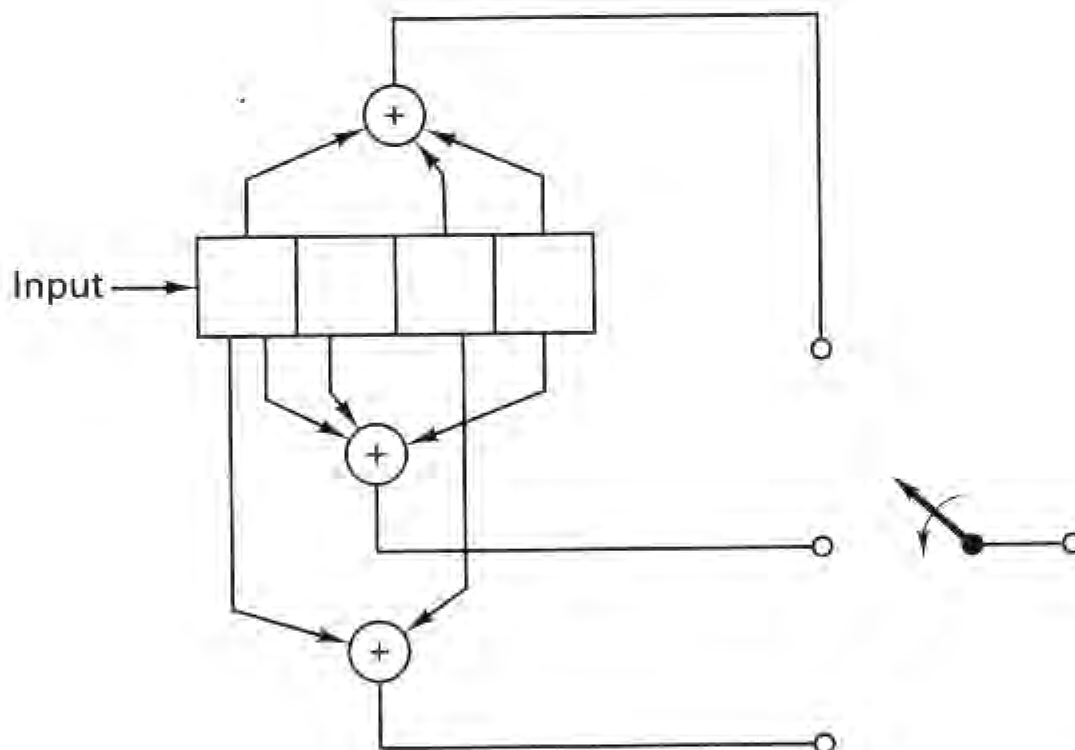

Figure P7.6

diagram to decode the first three data bits. Assume that the encoder had started in the 00 state, and that the decoding process is perfectly synchronized.

QUESTIONS

7.1. In convolutional encoding, why is *flushing* of the register periodically performed? (See Sections 7.2.1 and 7.3.4.)

7.2. Define what is meant by the *state* of a machine. (See Section 7.2.2.)

7.3. What is a *finite-state machine*? (See Section 7.2.2.)

7.4. What are *soft decisions,* and how much *greater complexity* is there in the process of soft-decision Viterbi decoding as compared with hard decision decoding? (See Sections 7.3.2 and 7.4.8.)

7.5. What is another (descriptive) name for a binary symmetric channel (BSC)? (See Section 7.3.2.1.)

7.6. Describe the *Add-Compare-Select* (ACS) computations performed in the process of Viterbi decoding. (See Section 7.3.5.)

7.7. On a trellis diagram, an *error* is associated with a surviving path that *diverges from,* and then *remerges to* the correct path. Why is it necessary for the path to remerge? (See Section 7.4.1.)

EXERCISES

Using the Companion CD, run the exercises associated with Chapter 7.