

LECTURE 6

Coping with Bit Errors

Recall our main goal in designing digital communication networks: to send information both reliably and efficiently between nodes. Meeting that goal requires the use of techniques to combat bit errors, which are an inevitable property of both communication channels and storage media.

The key idea we will apply to achieve reliable communication is *redundancy*: by replicating data to ensure that it can be reconstructed from some (correct) subset of the data received, we can improve the likelihood of fixing errors. This approach, called **error correction**, can work quite well when errors occur according to some random process. Error correction may not be able to correct all errors, however, so we will need a way to tell if any errors remain. That task is called **error detection**, which determines whether the data received after error correction is in fact the data that was sent. It will turn out that all the guarantees we can make are probabilistic, but we will be able to make our guarantees on reliable data receptions with very high probability. Over a communication channel, the sender and receiver implement the error correction and detection procedures. The sender has an *encoder* whose job is to take the message and process it to produce the *coded bits* that are then sent over the channel. The receiver has a *decoder* whose job is to take the received (coded) bits and to produce its best estimate of the message. The encoder-decoder procedures together constitute **channel coding**.

Our plan is as follows. First, we will define a model for the kinds of bit errors we're going to handle and revisit the previous lectures to see why errors occur. Then, we will discuss and analyze a simple redundancy scheme called a *replication code*, which will simply make c copies of any given bit. The replication code has a *code rate* of $1/c$ —that is, for every useful bit we receive, we will end up encoding c total bits. The overhead of the replication code of rate c is $1 - 1/c$, which is rather high for the error correcting power of the code. We will then turn to the key ideas in that allow us to build powerful codes capable of correcting errors without such a high overhead (or, capable of correcting far more errors at a given code rate than the trivial replication code).

There are two big ideas that are used in essentially all channel codes: the first is the notion of **embedding**, where the messages one wishes to send are placed in a geometrically pleasing way in a larger space so that the distance between any two valid points in the

embedding is large enough to enable the correction and detection of errors. The second big idea is to use **parity calculations** (or more generally, linear functions) over the bits we wish to send to produce the bits that are actually sent. We will study examples of embeddings and parity calculations in the context of two classes of codes: **linear block codes**, which are an instance of the broad class of **algebraic codes**, and **convolutional codes**, which are an instance of the broad class of **graphical codes**.¹

■ 6.1 Bit error models

In the previous lectures, we developed a model for how channels behave using the idea of linear time-invariance and saw how noise corrupted the information being received at the other end of a channel. We characterized the output of the channel, Y , as the sum of two components

$$y[n] = y_{\text{nf}}[n] + \text{noise}, \quad (6.1)$$

where $y[n]$ is the sum of two terms. The first term is the noise-free prediction of the channel output, which can be computed as the convolution of the channel's unit sample response with the input X , and the second is a random additive *noise* term. A good noise model for many real-world channels is Gaussian; such a model is has a special name: *additive white Gaussian noise*, or AWGN.² AWGN has mean 0 and is fully characterized by the variance, σ^2 . The larger the variance, the more intense the noise.

One of the properties of AWGN is that there is *always* a non-zero probability that a voltage transmitted to represent a 0 will arrive at the receiver with enough noise that it will be interpreted as a 1, and vice versa. For such a channel, if we know the transmitter's signaling levels and receiver's digitizing threshold, we know (from the earlier lecture on noise) how to calculate the *probability of bit error* (BER). For the most part, we will assume a simple (but useful) model that follows from the properties of AWGN: a transmitted 0 bit may be digitized at the receiver as a 1 (after deconvolution) with some probability p (the BER), and a transmitted 1 may be digitized as a 0 at the receiver (after deconvolution) with the same probability p . Such a model is also called a *binary symmetric channel*, or BSC. The "symmetric" refers to the property that a 0 becomes a 1 and vice versa with the same probability, p .

BSC is perhaps the simplest error model that is realistic, but real-world channels exhibit more complex behaviors. For example, over many wireless and wired channels as well as on storage media (like CDs, DVDs, and disks), errors can occur in *bursts*. That is, the probability of any given bit being received wrongly depends on (recent) history: the probability is higher if the bits in the recent past were received incorrectly.

Our goal is to develop techniques to mitigate the effects of both the BSC and burst errors. We'll start with techniques that work well over a BSC and then discuss how to deal with bursts.

A BSC error model is characterized by one number, p . We can determine p empirically by noting that if we send N bits over the channel, the expected number of erroneously

¹Graphical codes are sometimes also called "probabilistic codes" in the literature, for reasons we can't get into here.

²The "white" part of this term refers to the variance being the same over all the frequencies being used to communicate.

received bits is $N \cdot p$. Hence, by sending a known bit pattern and counting the fraction or erroneously received bits, we can estimate p . In practice, even when BSC is a reasonable error model, the range of p could be rather large, between 10^{-2} (or even a bit higher) all the way to 10^{-10} or even 10^{-12} . A value of p of about 10^{-2} means that messages longer than a 100 bits will see at least one error on average; given that the typical unit of communication over a channel (a “packet”) is generally at least 1000 bits long (and often bigger), such an error rate is too high.

But is a p of 10^{-12} small enough that we don’t need to bother about doing any error correction? The answer often depends on the data rate of the channel. If the channel has a rate of 10 Gigabits/s (available today even on commodity server-class computers), then the “low” p of 10^{-12} means that the receiver will see one error every 10 seconds on average if the channel is continuously loaded. Unless we take some mechanisms to mitigate the situation, the applications using the channel may find this error rate (in time) too frequent. On the other hand, a p of 10^{-12} may be quite fine over a communication channel running at 10 Megabits/s, as long as there is some way to detect errors when they occur.

It is important to note that the error rate is not a fixed property of the channel: it depends on the strength of the signal relative to the noise level (aka the “signal to noise ratio”, or SNR). Moreover, almost every communication channel trades-off between the raw transmission rate (“samples per bit” in this course) and the BER. As you increase the bit rate at which data is sent (say, by reducing the number of samples per bit), the BER will also increase, and then you need mechanisms to bring the BER down to acceptable values once more—and those mechanisms will reduce the achievable bit rate.

■ 6.2 The Simplest Code: Replication

In a replication code, each bit b is encoded as c copies of b , and the result is delivered. If we consider bit b to be the *message word*, then the corresponding *code word* is b^c (i.e., $bb\dots b$, c times). Clearly, there are only two possible message words (0 and 1) and corresponding code words. The beauty of the replication code is how absurdly simple it is.

But how well might it correct errors? To answer this question, we will write out the probability of decoding error for the BSC error model with the replication code. That is, if the channel corrupts each bit with probability p , what is the probability that the receiver decodes the received code word correctly to produce the message word that was sent?

The answer depends on the decoding method used. A reasonable decoding method is *maximum likelihood decoding*: given a received code word, r , which is some c -bit combination of 0’s and 1’s, the decoder should produce the most likely message that could have caused r to be received. If the BSC error probability is smaller than $1/2$, which will always be true for the realm in which coding should be used, then the most likely option is the message word that has the most number of bits in common with r .

Hence, the decoding process is as follows. First, count the number of 1’s in r . If there are more than $\frac{c}{2}$ 1’s, then decode the message as 1. If there are more than $\frac{c}{2}$ 0’s, then decode the message as 0. When c is odd, each code word will be decoded unambiguously. When c is even, and has an equal number of 0’s and 1’s, the decoder can’t really tell whether the message was a 0 or 1, and the best it can do is to make an arbitrary decision. (We have tacitly assumed that the *a priori* probability of the sender sending a message 0 is the same

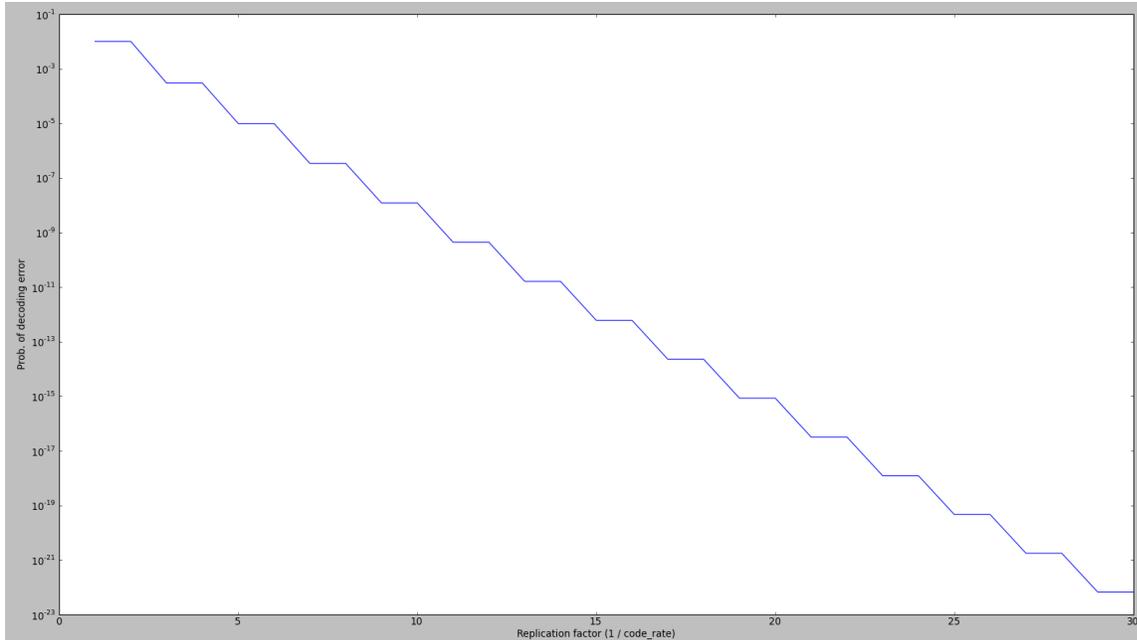


Figure 6-1: Probability of a decoding error with the replication code that replaces each bit b with c copies of b . The code rate is $1/c$.

as a 1.)

We can write the probability of decoding error for the replication code as, being a bit careful with the limits of the summation:

$$P(\text{decoding error}) = \begin{cases} \sum_{i=\lceil \frac{c}{2} \rceil}^c \binom{c}{i} p^i (1-p)^{c-i} & \text{if } c \text{ odd} \\ \sum_{i=\lceil \frac{c}{2} \rceil}^c \binom{c}{i} p^i (1-p)^{c-i} + \frac{1}{2} \binom{c}{c/2} p^{c/2} (1-p)^{c/2} & \text{if } c \text{ even} \end{cases} \quad (6.2)$$

When c is even, we add a term at the end to account for the fact that the decoder has a fifty-fifty chance of guessing correctly when it receives a codeword with an equal number of 0's and 1's.

Figure 6-1 shows the probability of decoding error from Eq.(6.2) as a function of the code rate for the replication code. The y -axis is on a log scale, and the probability of error is more or less a straight line with negative slope (if you ignore the flat pieces), which means that the decoding error probability decreases exponentially with the code rate. It is also worth noting that the error probability is the same when $c = 2\ell$ as when $c = 2\ell - 1$. The reason, of course, is that the decoder obtains no additional information that it already didn't know from any $2\ell - 1$ of the received bits.

Given a chunk of data of size s bits, we can now calculate the probability that it will be in error after the error correction code has been applied. Each message word (1 bit in the case of the replication code) will be decoded incorrectly with probability q , where q is given by Eq.(6.2). The probability that the entire chunk of data will be decoded correctly is given by $(1 - q)^s$, and the desired error probability is therefore equal to $1 - (1 - q)^s$. When $q \ll 1$, that error probability is approximately qs . This result should make intuitive sense.

Despite the exponential reduction in the probability of decoding error as c increases,

the replication code is extremely inefficient in terms of the overhead it incurs. As such, it is used only in situations when bandwidth is plentiful and there isn't much computation time to implement a more complex decoder.

We now turn to developing more sophisticated codes. There are two big ideas: *embedding messages into spaces in a way that achieves structural separation and parity (linear) computations over the message bits.*

■ 6.3 Embeddings and Hamming Distance

Let's start our investigation into error correction by examining the situations in which error detection and correction are possible. For simplicity, we will focus on single error correction (SEC) here.

There are 2^n possible n -bit strings. Let's define the *Hamming distance* (HD) between two n -bit words, w_1 and w_2 , as the number of bit positions in which the messages differ. Thus $0 \leq \text{HD}(w_1, w_2) \leq n$.

Suppose that $\text{HD}(w_1, w_2) = 1$. Consider what happens if we transmit w_1 and there's a single bit error that inconveniently occurs at the one bit position in which w_1 and w_2 differ. From the receiver's point of view it just received w_2 —the receiver can't detect the difference between receiving w_1 with a unfortunately placed bit error and receiving w_2 . In this case, we cannot guarantee that all single bit errors will be corrected if we choose a code where w_1 and w_2 are both valid code words.

What happens if we increase the Hamming distance between any two valid code words to at least 2? More formally, let's restrict ourselves to only sending some subset $\mathcal{S} = \{w_1, w_2, \dots, w_s\}$ of the 2^n possible words such that

$$\text{HD}(w_i, w_j) \geq 2 \text{ for all } w_i, w_j \in \mathcal{S} \text{ where } i \neq j \quad (6.3)$$

Thus if the transmission of w_i is corrupted by a single error, the result is *not* an element of \mathcal{S} and hence can be detected as an erroneous reception by the receiver, which knows which messages are elements of \mathcal{S} . A simple example is shown in Figure 6-2: 00 and 11 are valid code words, and the receptions 01 and 10 are surely erroneous.

It should be easy to see what happens as we use a code whose minimum Hamming distance between any two valid code words is D . We state the property formally:

Theorem 6.1 *A code with a minimum Hamming distance of D can detect any error pattern of $D - 1$ or fewer errors. Moreover, there is at least one error pattern with D errors that cannot be detected reliably.*

Hence, if our goal is to detect errors, we can use an embedding of the set of messages we wish to transmit into a bigger space, so that the minimum Hamming distance between any two code words in the bigger space is at least one more than the number of errors we wish to detect. (We will discuss how to produce such embeddings in the subsequent sections.)

But what about the problem of *correcting* errors? Let's go back to Figure 6-2, with $\mathcal{S} = \{00, 11\}$. Suppose the receiver receives 01. It can tell that a single error has occurred, but it can't tell whether the correct data sent was 00 or 11—both those possible patterns are equally likely under the BSC error model.

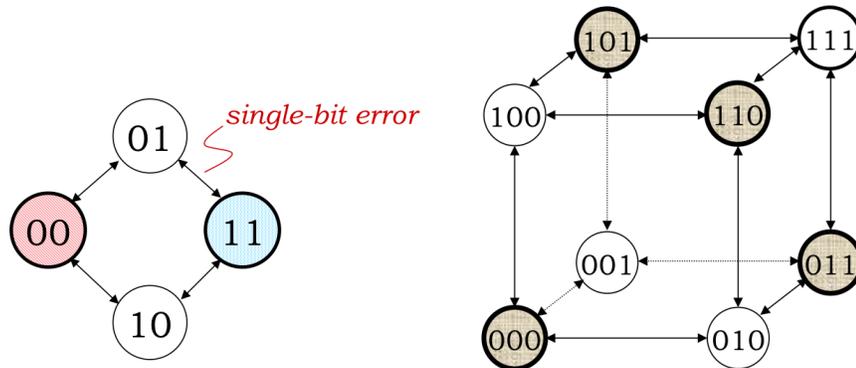


Figure 6-2: Code words separated by a Hamming distance of 2 can be used to detect single bit errors. The code words are shaded in each picture. The picture on the left is a (2,1) repetition code, which maps 1-bit messages to 2-bit code words. The code on the right is a (3,2) code, which maps 2-bit messages to 3-bit code words.

Ah, but we can extend our approach by producing an embedding with more space between valid codewords! Suppose we limit our selection of messages in \mathcal{S} even further, as follows:

$$\text{HD}(w_i, w_j) \geq 3 \text{ for all } w_i, w_j \in \mathcal{S} \text{ where } i \neq j \quad (6.4)$$

How does it help to increase the minimum Hamming distance to 3? Let's define one more piece of notation: let \mathcal{E}_{w_i} be the set of messages resulting from corrupting w_i with a single error. For example, $\mathcal{E}_{000} = \{001, 010, 100\}$. Note that $\text{HD}(w_i, \text{an element of } \mathcal{E}_{w_i}) = 1$.

With a minimum Hamming distance of 3 between the valid code words, observe that there is no intersection between \mathcal{E}_{w_i} and \mathcal{E}_{w_j} when $i \neq j$. Why is that? Suppose there was a message w_k that was in both \mathcal{E}_{w_i} and \mathcal{E}_{w_j} . We know that $\text{HD}(w_i, w_k) = 1$ and $\text{HD}(w_j, w_k) = 1$, which implies that w_i and w_j differ in at most two bits and consequently $\text{HD}(w_i, w_j) \leq 2$. That contradicts our specification that their minimum Hamming distance be 3. So the \mathcal{E}_{w_i} don't intersect.

Now we can correct single bit errors as well: the received message is either a member of \mathcal{S} (no errors), or is a member of some particular \mathcal{E}_{w_i} (one error), in which case the receiver can deduce the original message was w_i . Here's another simple example: let $\mathcal{S} = \{000, 111\}$. So $\mathcal{E}_{000} = \{001, 010, 100\}$ and $\mathcal{E}_{111} = \{110, 101, 011\}$ (note that \mathcal{E}_{000} doesn't intersect \mathcal{E}_{111}). Suppose the receiver receives 101. It can tell there's been a single error since $101 \notin \mathcal{S}$. Moreover it can deduce that the original message was 111 since $101 \in \mathcal{E}_{111}$.

We can formally state some properties from the above discussion, and state what the error-correcting power of a code whose minimum Hamming distance is at least D .

Theorem 6.2 *The Hamming distance between n -bit words satisfies the triangle inequality. That is, $\text{HD}(x, y) + \text{HD}(y, z) \geq \text{HD}(x, z)$.*

Theorem 6.3 *For a BSC error model with bit error probability $< 1/2$, the maximum likelihood decoding strategy is to map any received word to the valid code word with smallest Hamming distance from the received one (ties may be broken arbitrarily).*

Theorem 6.4 *A code with a minimum Hamming distance of D can correct any error pattern of*

$\lfloor \frac{D-1}{2} \rfloor$ or fewer errors. Moreover, there is at least one error pattern with $\lfloor \frac{D-1}{2} \rfloor + 1$ errors that cannot be corrected reliably.

Equation (6.4) gives us a way of determining if single-bit error correction can always be performed on a proposed set \mathcal{S} of transmission messages—we could write a program to compute the Hamming distance between all pairs of messages in \mathcal{S} and verify that the minimum Hamming distance was at least 3. We can also easily generalize this idea to check if a code can always correct more errors. And we can use the observations made above to decode any received word: just find the closest valid code word to the received one, and then use the known mapping between each distinct message and the code word to produce the message. That check may be exponential in the number of message bits we would like to send, but would be reasonable if the number of bits is small.

But how do we go about finding a good embedding (i.e., good code words)? This task isn't straightforward, as the following example shows. Suppose we want to reliably send 4-bit messages so that the receiver can correct all single-bit errors in the received words. Clearly, we need to find a set of messages \mathcal{S} with 2^4 elements. Quick, what should the members of \mathcal{S} be?

The answer isn't obvious. Once again, we could write a program to search through possible sets of n -bit messages until it finds a set of size 16 with a minimum Hamming distance of 3. An exhaustive search shows that the minimum n is 7, and one example of \mathcal{S} is:

0000000	1100001	1100110	0000111
0101010	1001011	1001100	0101101
1010010	0110011	0110100	1010101
1111000	0011001	0011110	1111111

But such exhaustive searches are impractical when we want to send even modestly longer messages. So we'd like some constructive technique for building \mathcal{S} . Much of the theory and practice of coding is devoted to finding such constructions and developing efficient encoding and decoding strategies.

Broadly speaking, there are two classes of code constructions, each with an enormous number of example instances. The first is the class of **algebraic block codes**. The second is the class of **graphical codes**. We will study two simple examples of **linear block codes**, which themselves are a sub-class of algebraic block codes: Hamming codes and rectangular parity codes. We also note that the replication code discussed in Section 6.2 is an example of a linear block code.

In later lectures, we will study **convolutional codes**, a sub-class of graphical codes.

■ 6.4 Linear Block Codes and Parity Calculations

Linear block codes are examples of algebraic block codes, which take the set of k -bit messages we wish to send (there are 2^k of them) and produce a set of 2^k code words, each n bits long ($n \geq k$) using *algebraic operations* over the block. The word “block” refers to the fact that any long bit stream can be broken up into k -bit blocks, which are then expanded to produce n -bit code words that are sent.

Such codes are also called (n, k) codes, where k message bits are combined to produce n code bits (so each code word has $n - k$ “redundancy” bits). Often, we use the notation (n, k, d) , where d refers to the minimum Hamming distance of the block code. The *rate* of a block code is defined as k/n ; the larger the rate, the less the overhead incurred by the code.

A linear code (whether a block code or not) produces code words from message bits by restricting the algebraic operations to *linear functions* over the message bits. By linear, we mean that any given bit in a valid code word is computed as the weighted sum of one or more original message bits. Linear codes, as we will see, are both powerful and efficient to implement. They are widely used in practice. In fact, *all* the codes we will study—including convolutional codes—are linear, as are most of the codes widely used in practice. We already looked at the properties of a simple linear block code: the replication code we discussed in Section 6.2 is a linear block code with parameters $(c, 1, c)$.

To develop a little bit of intuition about the linear operations, let’s start with a “hat” puzzle, which might at first seem unrelated to coding.

There are N people in a room, each wearing a hat colored red or blue, standing in a line in order of increasing height. Each person can see only the hats of the people in front, and does not know the color of his or her own hat. They play a game as a team, whose rules are simple. Each person gets to say one word: “red” or “blue”. If the word they say correctly guesses the color of their hat, the team gets 1 point; if they guess wrong, 0 points. Before the game begins, they can get together to agree on a protocol (i.e., what word they will say under what conditions). Once they determine the protocol, they stop talking, form the line, and are given their hats at random.

Can you think of a protocol that will maximize their score? What score does your protocol achieve?

A little bit of thought will show that there is a way to use the concept of **parity** to enable $N - 1$ of the people to correctly decode the colors of their hats. In general, the “parity” of a set of bits x_1, x_2, \dots, x_n is simply equal to $(x_1 + x_2 + \dots + x_n)$, where the addition is performed modulo 2 (it’s the same as taking the exclusive OR of the bits). Even parity occurs when the sum is 0 (i.e., the number of 1’s is even), while odd parity is when the sum is 1.

Parity, or equivalently, arithmetic modulo 2, has a special name: algebra in a *Galois Field* of order 2, also denoted \mathbb{F}_2 . A field must define rules for addition and multiplication. Addition in \mathbb{F}_2 is as stated above: $0 + 0 = 1 + 1 = 0$; $1 + 0 = 0 + 1 = 1$. Multiplication is as usual: $0 \cdot 0 = 0 \cdot 1 = 1 \cdot 0 = 0$; $1 \cdot 1 = 1$. Our focus in 6.02 will be on linear codes over \mathbb{F}_2 , but there are natural generalizations to fields of higher order (in particular, Reed Solomon codes, which are over Galois Fields of order 2^q).

A linear block code is characterized by the following rule (which is both a necessary and a sufficient condition for a code to be a linear block code):

Definition 6.1 *A block code is said to be linear if, and only if, the sum of any two code words is another code word.*

For example, the code defined by code words 000, 101, 011 is *not* a linear code, because $101 + 011 = 110$ is not a code word. But if we add 110 to the set, we get a linear code

because the sum of any two code words is another code word. The code 000, 101, 011, 110 has a minimum Hamming distance of 2 (that is, the smallest Hamming distance between any two code words in 2), and can be used to detect all single-bit errors that occur during the transmission of a code word. You can also verify that the minimum Hamming distance of this code is equal to the smallest number of 1's in a non-zero code word. In fact, that's a general property of all linear block codes, which we state formally below.

Theorem 6.5 *Define the weight of a code word as the number of 1's in the word. Then, the minimum Hamming distance of a linear block code is equal to the weight of the non-zero code word with the smallest weight.*

To see why, use the property that the sum of any two code words must also be a code word, and that the Hamming distance between any two code words is equal to the weight of their sum (i.e., $\text{weight}(u + v) = \text{HD}(u, v)$). We leave the complete proof of this theorem as a useful and instructive exercise for the reader.

The rest of this section shows how to construct linear block codes over \mathbb{F}_2 . For simplicity, and without much loss of generality, we will focus on correcting single-bit errors. We will show two ways of building the set \mathcal{S} of transmission messages such that the size of \mathcal{S} will allow us to send messages of some specific length, and to describe how the receiver can perform error correction on the (possibly corrupted) received messages. These are both examples of *single error correcting* (SEC) codes.

We will start with a simple *rectangular parity* code, then discuss the cleverer and more efficient *Hamming code* in Section 6.4.3.

■ 6.4.1 Rectangular Parity SEC Code

Let $\text{parity}(w)$ equal the sum over \mathbb{F}_2 of all the bits in word w . We'll use \cdot to indicate the concatenation (sequential joining) of two messages or a message and a bit. For any message (sequence of one or more bits), let $w = M \cdot \text{parity}(M)$. You should be able to confirm that $\text{parity}(w) = 0$. Parity lets us detect single errors because the set of code words w (each defined as $M \cdot \text{parity}(M)$) has a Hamming distance of 2.

If we transmit w when we want to send some message M , then the receiver can take the received word, r , and compute $\text{parity}(r)$ to determine if a single error has occurred. The receiver's parity calculation returns 1 if an odd number of the bits in the received message have been corrupted. When the receiver's parity calculation returns a 1, we say there has been a *parity error*.

This section describes a simple approach to building a SEC code by constructing multiple parity bits, each over various subsets of the message bits, and then using the resulting parity errors (or non-errors) to help pinpoint which bit was corrupted.

Rectangular code construction: Suppose we want to send a k -bit message M . Shape the k bits into a rectangular array with r rows and c columns, i.e., $k = rc$. For example, if $k = 8$, the array could be 2×4 or 4×2 (or even 8×1 or 1×8 , though those are a little less interesting). Label each data bit with subscript giving its row and column: the first bit would be d_{11} , the last bit d_{rc} . See Figure 6-3.

d_{11}	d_{12}	d_{13}	d_{14}	$p_row(M, 1)$
d_{21}	d_{22}	d_{23}	d_{24}	$p_row(M, 2)$
$p_col(M, 1)$	$p_col(M, 2)$	$p_col(M, 3)$	$p_col(M, 4)$	

Figure 6-3: A 2×4 arrangement for an 8-bit message with row and column parity.

0	1	1	0	0	1	0	0	1	1	0	1	1	1	1
1	1	0	1	1	0	0	1	0	1	1	1	1	0	1
1	0	1	1		1	0	1	0		1	0	0	0	
				(a)					(b)					(c)

Figure 6-4: Example received 8-bit messages. Which have an error?

Define $p_row(i)$ to be the parity of all the bits in row i of the array and let R be all the row parity bits collected into a sequence:

$$R = [p_row(1), p_row(2), \dots, p_row(r)]$$

Similarly, define $p_col(j)$ to be the parity of all the bits in column j of the array and let C be the all the column parity bits collected into a sequence:

$$C = [p_col(1), p_col(2), \dots, p_col(c)]$$

Figure 6-3 shows what we have in mind when $k = 8$.

Let $w = M \cdot R \cdot C$, i.e., the transmitted code word consists of the original message M , followed by the row parity bits R in row order, followed by the column parity bits C in column order. The length of w is $n = rc + r + c$. This code is linear because all the parity bits are linear functions of the message bits. The rate of the code is $rc/(rc + r + c)$.

We now prove that the rectangular parity code can correct all single-bit errors.

Proof of single-error correction property: This rectangular code is an SEC code for all values of r and c . We will show that it can correct all single bit errors by showing that its minimum Hamming distance is 3 (i.e., the Hamming distance between any two code words is at least 3). Consider two different uncoded messages, M_i and M_j . There are three cases to discuss:

- If M_i and M_j differ by a single bit, then the row and column parity calculations involving that bit will result in different values. Thus, the corresponding code words, w_i and w_j , will differ by three bits: the different data bit, the different row parity bit, and the different column parity bit. So in this case $HD(w_i, w_j) = 3$.
- If M_i and M_j differ by two bits, then either (1) the differing bits are in the same row, in which case the row parity calculation is unchanged but two column parity calculations will differ, (2) the differing bits are in the same column, in which case the column parity calculation is unchanged but two row parity calculations will differ, or (3) the differing bits are in different rows and columns, in which case there will be two row and two column parity calculations that differ. So in this case $HD(w_i, w_j) \geq$

4.

- If M_i and M_j differ by three or more bits, then in this case $\text{HD}(w_i, w_j) \geq 3$ because w_i and w_j contain M_i and M_j respectively.

Hence we can conclude that $\text{HD}(w_i, w_j) \geq 3$ and our simple “rectangular” code will be able to correct all single-bit errors.

Decoding the rectangular code: How can the receiver’s decoder correctly deduce M from the received w , which may or may not have a single bit error? (If w has more than one error, then the decoder does not have to produce a correct answer.)

Upon receiving a possibly corrupted w , the receiver checks the parity for the rows and columns by computing the sum of the appropriate data bits *and* the corresponding parity bit (all arithmetic in \mathbb{F}_2). This sum will be 1 if there is a parity error. Then:

- If there are no parity errors, then there has not been a single error, so the receiver can use the data bits as-is for M . This situation is shown in Figure 6-4(a).
- If there is single row or column parity error, then the corresponding parity bit is in error. But the data bits are okay and can be used as-is for M . This situation is shown in Figure 6-4(c), which has a parity error only in the fourth column.
- If there is one row and one column parity error, then the data bit in that row and column has an error. The decoder repairs the error by flipping that data bit and then uses the repaired data bits for M . This situation is shown in Figure 6-4(b), where there are parity errors in the first row and fourth column indicating that d_{14} should be flipped to be a 0.
- Other combinations of row and column parity errors indicate that multiple errors have occurred. There’s no “right” action the receiver can undertake because it doesn’t have sufficient information to determine which bits are in error. A common approach is to use the data bits as-is for M . If they happen to be in error, that will be detected when validating by the error detection method.

This recipe will produce the most likely message, M , from the received code word if there has been at most a single transmission error.

In the rectangular code the number of parity bits grows at least as fast as \sqrt{k} (it should be easy to verify that the smallest number of parity bits occurs when the number of rows, r , and the number of columns, c , are equal). Given a fixed amount of communication bandwidth, we’re interested in devoting as much of it as possible to sending message bits, not parity bits. Are there other SEC codes that have better code rates than our simple rectangular code? A natural question to ask is: *how little redundancy can we get away with and still manage to correct errors?*

The Hamming code uses a clever construction that uses the intuition developed while answering the question mentioned above. We answer this question next.

■ 6.4.2 How many parity bits are needed in a SEC code?

Let’s think about what we’re trying to accomplish with a SEC code: the goal is to correct transmissions with at most a single error. For a transmitted message of length n there are

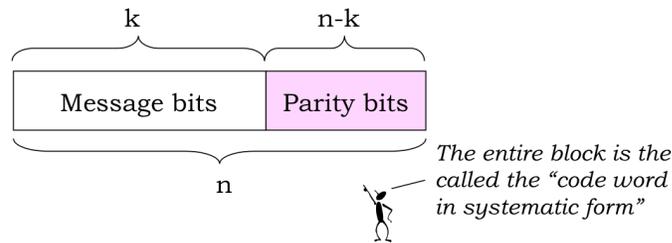


Figure 6-5: A code word in systematic form for a block code. Any linear code can be transformed into an equivalent systematic code.

$n + 1$ situations the receiver has to distinguish between: no errors and a single error in any of the n received bits. Then, depending on the detected situation, the receiver can make, if necessary, the appropriate correction.

Our first observation, which we will state here without proof, is that any linear code can be transformed into a **systematic** code. A systematic code is one where every n -bit code word can be represented as the original k -bit message followed by the $n - k$ parity bits (it actually doesn't matter how the original message bits and parity bits are interspersed). Figure 6-5 shows a code word in systematic form.

So, given a systematic code, how many parity bits do we absolutely need? We need to choose n so that single error correction is possible. Since there are $n - k$ parity bits, each combination of these bits must represent *some* error condition that we must be able to correct (or infer that there were no errors). There are 2^{n-k} possible distinct parity bit combinations, which means that we can distinguish at most that many error conditions. We therefore arrive at the constraint

$$n + 1 \leq 2^{n-k} \quad (6.5)$$

i.e., there have to be enough parity bits to distinguish all corrective actions that might need to be taken. Given k , we can determine the number of parity bits ($n - k$) needed to satisfy this constraint. Taking the log base 2 of both sides, we can see that the number of parity bits **must** grow at least *logarithmically* with the number of message bits. Not all codes achieve this minimum (e.g., the rectangular code doesn't), but the Hamming code, which we describe next, does.

■ 6.4.3 Hamming Codes

Intuitively, it makes sense that for a code to be efficient, each parity bit should protect as many data bits as possible. By symmetry, we'd expect each parity bit to do the same amount of "work" in the sense that each parity bit would protect the same number of data bits. If some parity bit is shirking its duties, it's likely we'll need a larger number of parity bits in order to ensure that each possible single error will produce a unique combination of parity errors (it's the unique combinations that the receiver uses to deduce which bit, if any, had a single error).

The class of Hamming single error correcting codes is noteworthy because they are particularly efficient in the use of parity bits: the number of parity bits used by Hamming

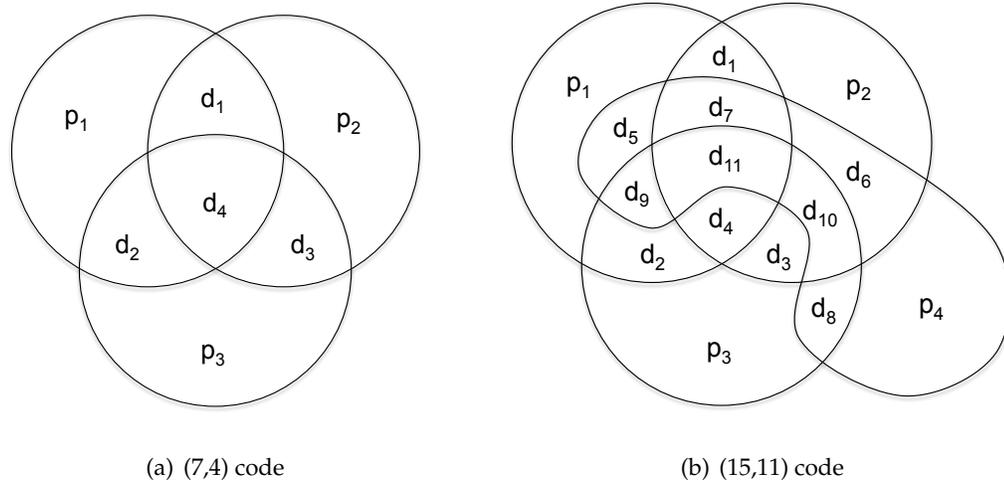


Figure 6-6: Venn diagrams of Hamming codes showing which data bits are protected by each parity bit.

codes grows logarithmically with the size of the code word.

Figure 6-6 shows two examples of the class: the (7,4) and (15,11) Hamming codes. The (7,4) Hamming code uses 3 parity bits to protect 4 data bits; 3 of the 4 data bits are involved in each parity computation. The (15,11) Hamming code uses 4 parity bits to protect 11 data bits, and 7 of the 11 data bits are used in each parity computation (these properties will become apparent when we discuss the logic behind the construction of the Hamming code in Section 6.4.4).

Looking at the diagrams, which show the data bits involved in each parity computation, you should convince yourself that each possible single error (don't forget errors in one of the parity bits!) results in a unique combination of parity errors. Let's work through the argument for the (7,4) Hamming code. Here are the parity-check computations performed by the receiver:

$$\begin{aligned} E_1 &= (d_1 + d_2 + d_4 + p_1) \pmod 2 \\ E_2 &= (d_1 + d_3 + d_4 + p_2) \pmod 2 \\ E_3 &= (d_2 + d_3 + d_4 + p_3) \pmod 2 \end{aligned}$$

where each E_i is called a *syndrome* bit because it helps the receiver diagnose the "illness" (errors) in the received data. For each combination of syndrome bits, we can look for the bits in each code word that appear in *all* the E_i computations that produced 1; these bits are potential candidates for having an error since any of them could have caused the observed parity errors. Now eliminate from the candidates bits that appear in *any* E_i computations that produced 0 since those calculations prove those bits didn't have errors. We'll be left with either no bits (no errors occurred) or one bit (the bit with the single error).

For example, if $E_1 = 1$, $E_2 = 0$ and $E_3 = 1$, we notice that bits d_2 and d_4 both appear in the computations for E_1 and E_3 . However, d_4 appears in the computation for E_2 and should be eliminated, leaving d_2 as the sole candidate as the bit with the error.

Another example: suppose $E_1 = 1$, $E_2 = 0$ and $E_3 = 0$. Any of the bits appearing in the

computation for E_1 could have caused the observed parity error. Eliminating those that appear in the computations for E_2 and E_3 , we're left with p_1 , which must be the bit with the error.

Applying this reasoning to each possible combination of parity errors, we can make a table that shows the appropriate corrective action for each combination of the syndrome bits:

$E_3E_2E_1$	Corrective Action
000	no errors
001	p_1 has an error, flip to correct
010	p_2 has an error, flip to correct
011	d_1 has an error, flip to correct
100	p_3 has an error, flip to correct
101	d_2 has an error, flip to correct
110	d_3 has an error, flip to correct
111	d_4 has an error, flip to correct

■ 6.4.4 Is There a Logic to the Hamming Code Construction?

So far so good, but the allocation of data bits to parity-bit computations may seem rather arbitrary and it's not clear how to build the corrective action table except by inspection.

The cleverness of Hamming codes is revealed if we order the data and parity bits in a certain way and assign each bit an index, starting with 1:

index	1	2	3	4	5	6	7
binary index	001	010	011	100	101	110	111
(7,4) code	p_1	p_2	d_1	p_3	d_2	d_3	d_4

This table was constructed by first allocating the parity bits to indices that are powers of two (e.g., 1, 2, 4, ...). Then the data bits are allocated to the so-far unassigned indices, starting with the smallest index. It's easy to see how to extend this construction to any number of data bits, remembering to add additional parity bits at indices that are a power of two.

Allocating the data bits to parity computations is accomplished by looking at their respective indices in the table above. Note that we're talking about the *index* in the table, not the subscript of the bit. Specifically, d_i is included in the computation of p_j if (and only if) the logical AND of $\text{index}(d_i)$ and $\text{index}(p_j)$ is non-zero. Put another way, d_i is included in the computation of p_j if, and only if, $\text{index}(p_j)$ contributes to $\text{index}(d_i)$ when writing the latter as sums of powers of 2.

So the computation of p_1 (with an index of 1) includes all data bits with odd indices: d_1 , d_2 and d_4 . And the computation of p_2 (with an index of 2) includes d_1 , d_3 and d_4 . Finally, the computation of p_3 (with an index of 4) includes d_2 , d_3 and d_4 . You should verify that these calculations match the E_i equations given above.

If the parity/syndrome computations are constructed this way, it turns out that $E_3E_2E_1$, treated as a binary number, gives the index of the bit that should be corrected. For example, if $E_3E_2E_1 = 101$, then we should correct the message bit with index 5, i.e., d_2 . This corrective action is exactly the one described in the earlier table we built by inspection.

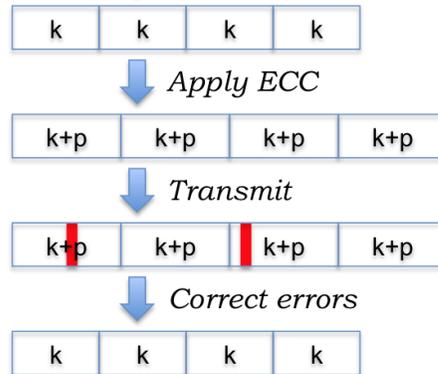


Figure 6-7: Dividing a long message into multiple SEC-protected blocks of k bits each, adding parity bits to each constituent block. The red vertical rectangles refer are bit errors.

The Hamming code's syndrome calculation and subsequent corrective action can be efficiently implemented using digital logic and so these codes are widely used in contexts where single error correction needs to be fast, e.g., correction of memory errors when fetching data from DRAM.

■ 6.5 Protecting Longer Messages with SEC Codes

SEC codes are a good building block, but they correct at most one error. As messages get longer, they are unlikely to provide much correction if we use the entire message as a single block of k bits. The solution, of course, is to break up a longer message into smaller blocks of k bits each, and to protect each one with its own SEC code. The result might look as shown in Figure 6-7.

■ 6.5.1 Coping with Burst Errors

Over many channels, errors occur in bursts and the BSC error model is invalid. For example, wireless channels suffer from *interference* from other transmitters and from *fading*, caused mainly by *multi-path propagation* when a given signal arrives at the receiver from multiple paths and interferes in complex ways because the different copies of the signal experience different degrees of attenuation and different delays. Another reason for fading is the presence of obstacles on the path between sender and receiver; such fading is called *shadow fading*.

The behavior of a fading channel is complicated and beyond the scope of 6.02, but the impact of fading on communication is that the random process describing the bit error probability is no longer independent and identically distributed. The BSC model needs to be replaced with a more complicated one in which errors may occur in *bursts*. Many such theoretical models guided by empirical data exist, but we won't go into them here. Our goal is to understand how to develop error correction mechanisms when errors occur in bursts.

But what do we mean by a "burst"? The simplest model is to model the channel as having two states, a "good" state and a "bad" state. In the "good" state, the bit error

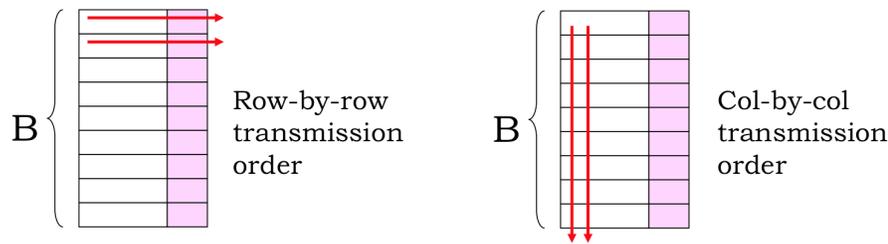


Figure 6-8: Interleaving can help recover from burst errors: code each block row-wise with an SEC, but transmit them in interleaved fashion in columnar order. As long as a set of burst errors corrupts some set of k^{th} bits, the receiver can recover from *all* the errors in the burst.

probability is p_g and in the “bad” state, it is p_b ; $p_b > p_g$. Once in the good state, the channel has some probability of remaining there (generally $> 1/2$) and some probability of moving into the “bad” state, and vice versa. It should be easy to see that this simple model has the property that the probability of a bit error depends on whether the previous bit (or previous few bits) are in error or not. The reason is that the odds of being in a “good” state are high if the previous few bits have been correct.

At first sight, it might seem like the SEC schemes we studied are poorly suited for a channel experiencing burst errors. The reason is shown in Figure 6-8 (left), where each block of the message is protected by its SEC parity bits. The different blocks are shown as different rows. When a burst error occurs, multiple bits in an SEC block are corrupted, and the SEC can’t recover from them.

Interleaving is a commonly used technique to recover from burst errors on a channel even when the individual blocks are protected with a code that, on the face of it, is not suited for burst errors. The idea is simple: code the blocks as before, but transmit them in a “columnar” fashion, as shown in Figure 6-8 (right). That is, send the first bit of block 1, then the first bit of block 2, and so on until all the first bits of each block are sent. Then, send the second bits of each block in sequence, then the third bits, and so on.

What happens on a burst error? Chances are that it corrupts a set of “first” bits, or a set of “second” bits, or a set of “third” bits, etc., because those are the bits sent in order on the channel. As long as only a set of k^{th} bits are corrupted, the receiver can correct *all* the errors. The reason is that each coded block will now have at most one error. Thus, SEC codes are a useful primitive to correct against burst errors, in concert with interleaving.

■ Acknowledgments

Many thanks to Katrina LaCurts for carefully reading these notes and making several useful comments.

■ Problems and Questions

These questions are to help you improve your understanding of the concepts discussed in this lecture. The ones marked ***PSet*** are in the online problem set or lab.

1. Show that the Hamming distance satisfies the triangle inequality. That is, show that $\text{HD}(x, y) + \text{HD}(y, z) \geq \text{HD}(x, z)$ for any three n -bit binary numbers in \mathbb{F}_2 .
2. Consider the following rectangular linear block code:

D0	D1	D2	D3	D4		P0
D5	D6	D7	D8	D9		P1
D10	D11	D12	D13	D14		P2

P3	P4	P5	P6	P7		

Here, D_0 – D_{14} are data bits, P_0 – P_2 are row parity bits and P_3 – P_7 are column parity bits. What are n , k , and d for this linear code?

3. ***PSet*** Consider a rectangular parity code as described in Section 6.4.1. Ben Bitdiddle would like use this code at a variety of different code rates and experiment with them on some channel.
 - (a) Is it possible to obtain a rate lower than $1/3$ with this code? Explain your answer.
 - (b) Suppose he is interested in code rates like $1/2$, $2/3$, $3/4$, etc.; i.e., in general a rate of $\frac{n-1}{n}$, for integer $n > 1$. Is it always possible to pick the parameters of the code (i.e., the block size and the number of rows and columns over which to construct the parity bits) so that any such code rate is achievable? Explain your answer.
4. Two-Bit Communications (TBC), a slightly suspect network provider, uses the following linear block code over its channels. All arithmetic is in \mathbb{F}_2 .

$$P_0 = D_0, P_1 = (D_0 + D_1), P_2 = D_1.$$

- (a) What are n and k for this code?
- (b) Suppose we want to perform syndrome decoding over the received bits. Write out the three syndrome equations for E_0, E_1, E_2 .
- (c) For the eight possible syndrome values, determine what error can be detected (none, error in a particular data or parity bit, or multiple errors). Make your choice using maximum likelihood decoding, assuming a small bit error probability (i.e., the smallest number of errors that's consistent with the given syndrome).
- (d) Suppose that the the 5-bit blocks arrive at the receiver in the following order: D_0, D_1, P_0, P_1, P_2 . If 11011 arrives, what will the TBC receiver report as the received data after error correction has been performed? Explain your answer.

(e) TBC would like to improve the code rate while still maintaining single-bit error correction. Their engineer would like to reduce the number of parity bits by 1. Give the formulas for P_0 and P_1 that will accomplish this goal, or briefly explain why no such code is possible.

5. For any linear block code over \mathbb{F}_2 with minimum Hamming distance at least $2t + 1$ between code words, show that:

$$2^{n-k} \geq 1 + \binom{n}{1} + \binom{n}{2} + \dots + \binom{n}{t}.$$

Hint: How many errors can such a code always correct?

6. Using the Hamming code construction for the (7,4) code, construct the parity equations for the (15,11) code. How many equations does this code have? How many message bits contribute to each parity bit?
7. Prove Theorems 6.2 and 6.3. (Don't worry too much if you can't prove the latter; we will give the proof when we discuss convolutional codes in Lecture 8.)
8. The weight of a code word in a linear block code over \mathbb{F}_2 is the number of 1's in the word. Show that any linear block code must either: (1) have only even weight code words, or (2) have an equal number of even and odd weight code words.

Hint: Proof by contradiction.