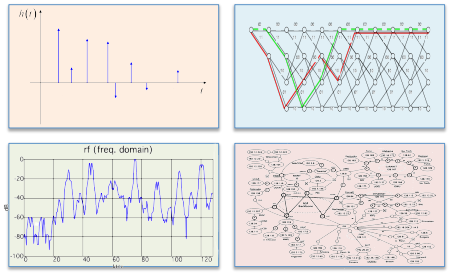

INTRODUCTION TO EECS II

DIGITAL COMMUNICATION SYSTEMS

6.02 Spring 2010
Lecture #6

- Coping with errors using packets
- Checksums, CRC
- Hamming distance & single error correction
- (n,k) block codes

There's good news and bad news...

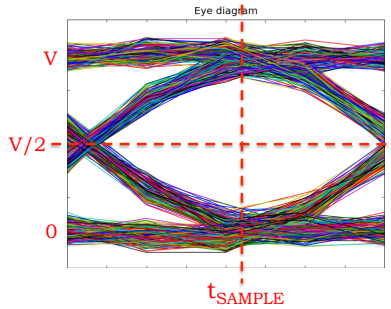

The good news: Our digital signaling scheme usually allows us to recover the original signal despite small amplitude errors introduced by inter-symbol interference and noise. An example of the digital abstraction doing its job!

The bad news: larger amplitude errors (hopefully infrequent) that change the signal irretrievably. These show up as bit errors in our digital data stream.

Bit Errors

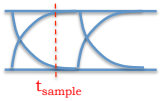

Assuming a Gaussian PDF for noise and only 1-bit of inter-symbol interference, samples at t_{SAMPLE} have the following PDF:

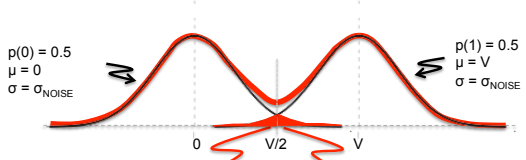

p(1)*p(rcv 0 | xmit 1) p(0)*p(rcv 1 | xmit 0)

We can estimate the bit-error rate (BER) using Φ, the unit normal cumulative distribution function:

$$BER = (0.5)\Phi\left[\frac{V/2 - V}{\sigma_{NOISE}}\right] + (0.5)\left[1 - \Phi\left[\frac{V/2 - 0}{\sigma_{NOISE}}\right]\right] = \Phi\left[\frac{-V/2}{\sigma_{NOISE}}\right]$$

For a smaller BER, you need a smaller σ_{NOISE} or a larger V!

Dealing With Errors: Packets

To deal with errors, divide message into fixed-sized packets, which are transmitted one after another.

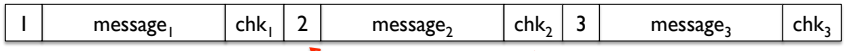

Packet = {#, message, chk}

Sequence number provides unique identifier for each packet.

Checksum = f(#,message)
Receiver recomputes checksum and compares with received checksum: mismatch → packet error, so request retransmission using sequence number.

Checking a Packet for Errors

- Transmitter appends check bytes computed from sequence number and data to end of packet

- Computing the check bytes
  - Checksum – easy in SW
  - Cyclical Redundancy Check (CRC) – easy in HW

- Checksum:
  - Add up all the message units and send along sum
  - Adler-32: two 16-bit mod-65521 sums A and B
    - A is 1 + sum of all the bytes in the message, mod 65521
    - B is the sum of the A values after each addition, mod 65521
    - 32-bit checksum = (B<<16)+A

Cyclical Redundancy Check

Example: CRC-16

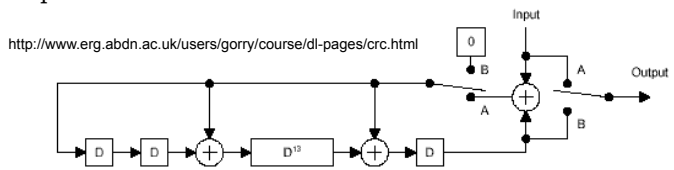

Sending: Initialize all D elements to 0. Set switch to position A, send message bit-by-bit. When complete, set switch to position B and send 16 more bits.

Receiving: Initialize all D elements to 0. Set switch to position A, receive message and CRC bit-by-bit. If correct, all D elements should be 0 after last bit has been processed.

CRC-16 detects all single- and double-bit errors, all odd numbers of errors, all errors with burst lengths < 16, and 99.984% of all other errors.

BER for common channels

Channel type	Bandwidth	BER
Telephone Landline	2 Mbits/sec	10^{-4} to 10^{-6}
Twisted pair (differential)	1 Gbits/sec	$\leq 10^{-7}$
Coaxial cable	100 Mbits/sec	$\leq 10^{-6}$
Fiber Optics	10 Tbits/sec	$\leq 10^{-9}$
Infrared	2 Mbits/sec	10^{-4} to 10^{-6}
3G cellular	1 Mbits/sec	10^{-4}

Source: Rahmani, et al, *Error Detection Capabilities of Automotive Technologies and Ethernet – A Comparative Study*, 2007 IEEE Intelligent Vehicles Symposium, p 674-679

How Frequent is Packet Retransmission?

p(1 or more errors) = 1 – p(no errors) = $1 - (1 - \text{BER})^k$

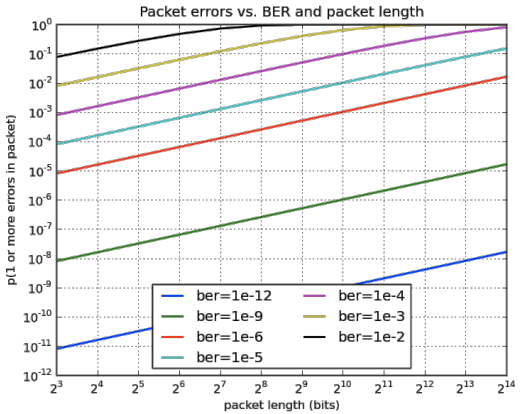

With 1kbyte packets and BER=1e-6, retransmit 1 every 100.

Implement Single Error Correction?

To reduce retransmission rate, suppose we invent a scheme that can correct single-bit errors and apply it to sub-blocks of the data packet (effectively reducing k). Does that help?

p(2 or more errors) = 1 – p(no errors) – p(exactly one error)
= 1 – $(1 - BER)^k$ – k*BER*$(1-BER)^{k-1}$

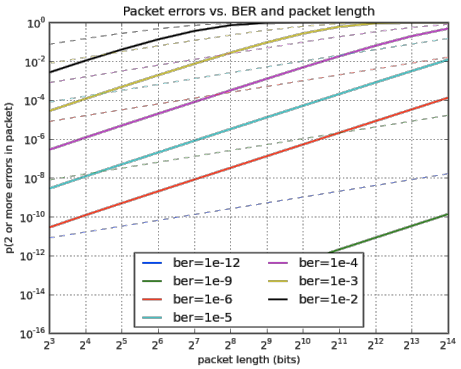

Digital Transmission using SECC

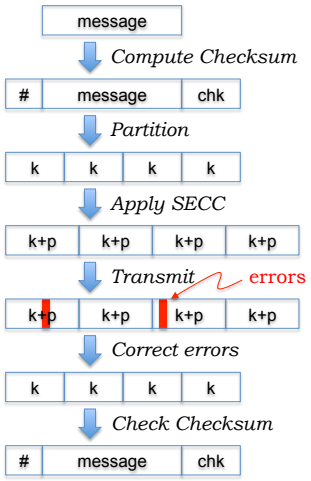

- Start with original message
- Add checksum to enable verification of error-free transmission
- Apply SECC, adding parity bits to each k-bit block of the message. Number of parity bits (p) depends on code:
  - Replication: p grows as O(k)
  - Rectangular: p grows as O(√k)
  - Hamming: p grows as O(log k)
- After xmit, correct errors
- Verify checksum, fails if undetected/uncorrectable error
- Deliver or discard message

Channel coding

Our plan to deal with bit errors:

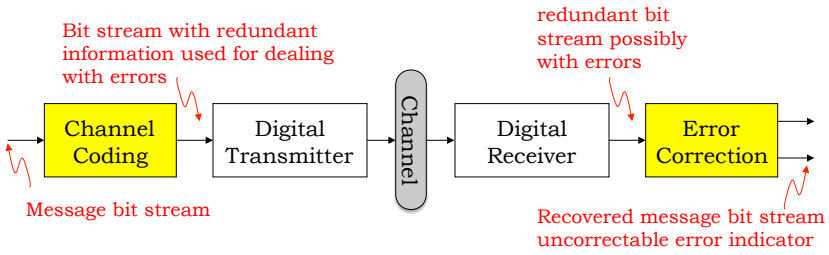

We'll add redundant information to the transmitted bit stream (a process called channel coding) so that we can detect errors at the receiver. Ideally we'd like to correct commonly occurring errors, e.g., error bursts of bounded length. Otherwise, we should detect uncorrectable errors and use, say, retransmission to deal with the problem.

Error detection and correction

Suppose we wanted to reliably transmit the result of a single coin flip:


Heads: "0"

Tails: "1"

This is a prototype of the "bit" coin for the new information economy. Value = 12.5¢

Further suppose that during transmission a single-bit error occurs, i.e., a single "0" is turned into a "1" or a "1" is turned into a "0".

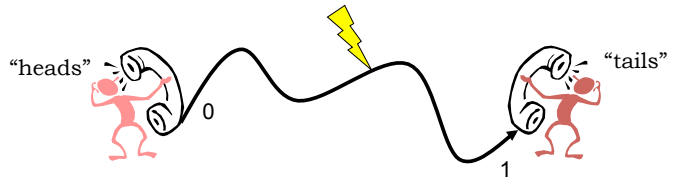

Hamming Distance

(Richard Hamming, 1950)

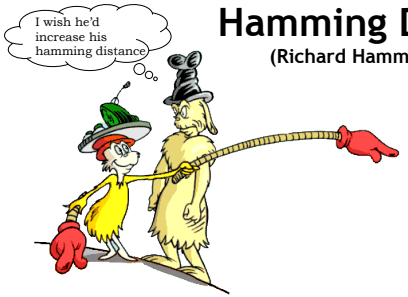

HAMMING DISTANCE: The number of digit positions in which the corresponding digits of two encodings of the same length are different

The Hamming distance between a valid binary code word and the same code word with single-bit error is 1.

The problem with our simple encoding is that the two valid code words ("0" and "1") also have a Hamming distance of 1. So a single error changes a valid code word into another valid code word…

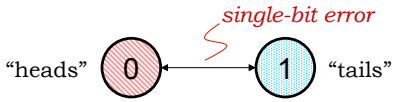

Error Detection

What we need is an encoding where a single-bit error doesn't produce another valid code word.

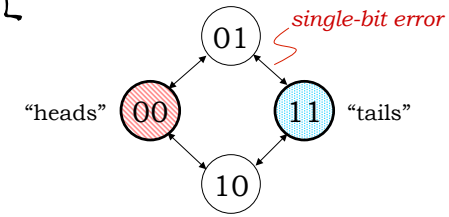

If D is the minimum Hamming distance between code words, we can detect up to (D-1)-bit errors

We can add single error detection to any length code word by adding a *parity bit* chosen to guarantee the Hamming distance between any two valid code words is at least 2. In the diagram above, we're using "even parity" where the added bit is chosen to make the total number of 1's in the code word even.

Parity check

- A parity bit can be added to any length message and is chosen to make the total number of "1" bits even (aka "even parity").
- To check for a single-bit error (actually any odd number of errors), count the number of "1"s in the received message and if it's odd, there's been an error.

0 1 1 0 0 1 0 1 0 0 1 1 → original word with parity
0 1 1 0 0 0 0 1 0 0 1 1 → single-bit error (detected)
0 1 1 0 0 0 1 1 0 0 1 1 → 2-bit error (not detected)

- One can "count" by summing the bits in the word modulo 2 (which is equivalent to XOR'ing the bits together).

Error Correction

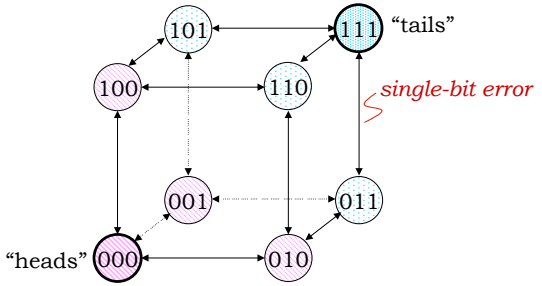

If D is the minimum Hamming distance between code words, we can correct up to $\left\lfloor \frac{D-1}{2} \right\rfloor$ bit errors

By increasing the Hamming distance between valid code words to 3, we guarantee that the sets of words produced by single-bit errors don't overlap. So if we detect an error, we can perform *error correction* since we can tell what the valid code was before the error happened.

- Can we safely detect double-bit errors while correcting 1-bit errors?
- Do we always need to triple the number of bits?

Single Error Correcting Codes (SECC)

Basic idea:

- Use multiple parity bits, each covering a subset of the data bits.
- No two message bits belong to exactly the same subsets, so a single error will generate a unique set of parity check errors.

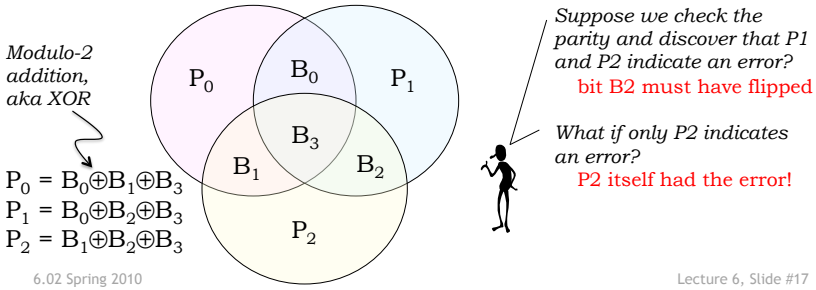

Checking the parity

- Transmit: Compute the parity bits and send them along with the message bits
- Receive: After receiving the (possibly corrupted) message, compute a syndrome bit (E_i) for each parity bit. For the code on previous slide:

$$E_0 = B_0 \oplus B_1 \oplus B_3 \oplus P_0$$
$$E_1 = B_0 \oplus B_2 \oplus B_3 \oplus P_1$$
$$E_2 = B_1 \oplus B_2 \oplus B_3 \oplus P_2$$

- If all the E_i are zero: no errors!
- Otherwise the particular combination of the E_i can be used to figure out which bit to correct.

Using the Syndrome to Correct Errors

Continuing example from previous slides: there are three syndrome bits, giving us a total of 8 encodings.

$E_2E_1E_0$	Single Error Correction
0 0 0	No errors
0 0 1	P0 has an error, flip to correct
0 1 0	P1 has an error, flip to correct
0 1 1	B0 has an error, flip to correct
1 0 0	P2 has an error, flip to correct
1 0 1	B1 has an error, flip to correct
1 1 0	B2 has an error, flip to correct
1 1 1	B3 has an error, flip to correct

What happens if there is more than one error?

The 8 encodings indicate the 8 possible correction actions: no errors, error in one of 4 data bits, error in one of 3 parity bits

(n,k,d) Systematic Block Codes

- Split message into *k*-bit blocks
- Add (*n-k*) parity bits to each block, making each block *n* bits long.

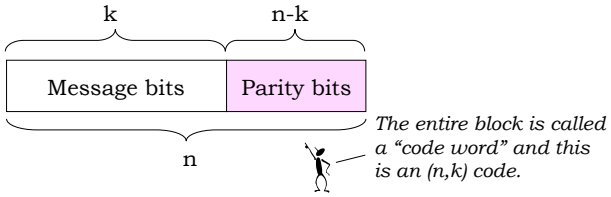

- Often we'll use the notation (n,k,d) where d is the minimum Hamming distance between code words.
- The ratio k/n is called the *code rate* and is a measure of the code's overhead (always ≤ 1, larger is better).

A simple (8,4,3) code

Idea: start with rectangular array of data bits, add parity checks for each row and column. Single-bit error in data will show up as parity errors in a particular row and column, pinpointing the bit that has the error.

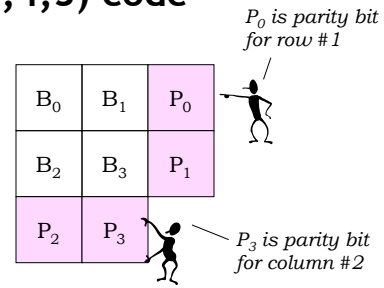

0 1 1
1 1 0
1 0

Parity for each row and column is correct ⇒ no errors

0 1 1
1 0 0
1 0

Parity check fails for row #2 and column #2 ⇒ bit B_3 is incorrect

0 1 1
1 1 1
1 0

Parity check only fails for row #2 ⇒ bit P_1 is incorrect

Can you verify this code has a Hamming distance of 3?

Double-error Detection

Suppose we wanted to detect 2-bit errors as well as correct 1-bit errors. Then we'd need a code with a Hamming distance of 4 instead of 3.

For example, we can generate a (9,4,4) code from our (8,4,3) code by adding an overall parity bit:

$$P_4 = B_0 \oplus B_1 \oplus B_2 \oplus B_3 \oplus P_0 \oplus P_1 \oplus P_2 \oplus P_3$$

P_4 indicates if there have been an odd number of errors in the other data and parity bits, which is enough to enumerate all the possibilities of 0, 1 or 2 errors:

- No error (all P_i are 0)
- Correctable 1-bit error
  - $P_4 = 1 \rightarrow$ odd number of errors $\rightarrow$ 1 error
- Uncorrectable 2-bit errors
  - $P_4 = 0$, other P_i nonzero $\rightarrow$ there are an even number of errors $\rightarrow$ 2 errors

Summary

- Divide message into separate transmission units called packets
  - Each packet has a sequence number, checksum
  - Checksum used to verify if packet has been received correctly. If not, request retransmission
  - Sequence number used in retransmission requests, detecting lost packets
- Single error correction greatly reduces number of packets that would need to be retransmitted
  - Convert k-bit sub-block of packet to n-bit code word
  - Require n-bit code words to have min Hamming distance of 3 $\rightarrow$ single error correction possible