

CSE 461: Framing, Error Detection and Correction

Next Topics

- Framing
 - Focus: How does a receiver know where a message begins/ends
- Error detection and correction
 - Focus: How do we detect and correct messages that are garbled during transmission?
 - The responsibility for doing this cuts across the different layers




Framing

- Need to send message, not just bits
 - Requires that we synchronize on the start of message reception at the far end of the link
 - Complete Link layer messages are called frames
- Common approaches: Sentinels, lengths, clock-based
 - Sentinels: Look for special control code that marks start of frame and escape or “stuff” this code within the data region
 - Lengths: additionally, tell receiver how large the frame is going to be
 - Clocks: agree on when frames ought to begin/end

Example: Point-to-Point Protocol (PPP)

- IETF standard, used for dialup and leased lines



- Flag is special and indicates start/end of frame
- Occurrences of flag inside payload must be "stuffed"
 - Like an "escape" character:
 - " → \"
 - Replace 0x7E with 0x7D, 0x5E
 - Replace 0x7D with 0x7D, 0x5D
- Problems?
- Why not use a length field?

Example: SONET

- Standard for long distance transmission over optical networks
 - Base rate STS-1 of 51.84 Mbps
 - STS-192 in use
- All packets are 125 μ s long \rightarrow start frame synch bytes just used for synchronization
- No character stuffing, no need for length field to find end of frames
 - Why not always do things SONET style?

Error Detection and Redundancy

- Noise can flip some of the bits we receive
 - We must be able to detect when this occurs!
 - Who needs to detect it? (links/routers, OSs, or apps?)
- Basic approach: add redundant data
 - Error detection codes allow errors to be recognized
 - Error correction codes allow errors to be repaired too

Motivating Example

- A simple error detection scheme:
 - Just send two copies. Differences imply errors.
- Question: Can we do any better?
 - With less overhead
 - Catch more kinds of errors
- Answer: Yes – stronger protection with fewer bits
 - But we can't catch all inadvertent errors, nor malicious ones
- We will look at basic block codes
 - K bits in, N bits out is a (N,K) code
 - Simple, memoryless mapping

The Hamming Distance


- Errors must not turn one valid codeword into another valid codeword, or we cannot detect/correct them.
- Hamming distance of a code is the smallest number of bit differences that turn any one codeword into another
 - e.g, code 000 for 0, 111 for 1, Hamming distance is 3
- For code with distance $d+1$:
 - d errors can be detected, e.g, 001, 010, 110, 101, 011
- For code with distance $2d+1$:
 - d errors can be corrected, e.g., 001 \rightarrow 000

Parity

- Start with n bits and add another so that the total number of 1s is even (even parity)
 - e.g. 0110010 → 01100101
 - Easy to compute as XOR of all input bits
- Will detect an odd number of bit errors
 - But not an even number
- Does not correct any errors

2D Parity

- Add parity row/column to array of bits
- How many simultaneous bit errors can it detect?
- Which errors can it correct?



Checksums


- Used in Internet protocols (IP, ICMP, TCP, UDP)
- Basic Idea: Add up the data and send it along with sum
- Algorithm:
 - checksum is the 1s complement of the 1s complement sum of the data interpreted 16 bits at a time (for 16-bit TCP/UDP checksum)
- 1s complement: flip all bits to make number negative
 - Consequence: adding requires carryout to be added back

CRCs (Cyclic Redundancy Check)

- Stronger protection than checksums
 - Used widely in practice, e.g., Ethernet CRC-32
 - Implemented in hardware (XORs and shifts)
- Algorithm: Given n bits of data, generate a k bit check sequence that gives a combined $n + k$ bits that are divisible by a chosen divisor $C(x)$
- Based on mathematics of finite fields
 - “numbers” correspond to polynomials, use modulo arithmetic
 - e.g, interpret 10011010 as $x^7 + x^4 + x^3 + x^1$

CRC Example

- Extend message with k 0's, when using a k -degree generator
- Divide message by generator (XOR)
- Discard result
- Subtract remainder from original message
- On reception, check that message is divisible by generator



How is $C(x)$ Chosen?

- Mathematical properties:
 - All 1-bit errors if non-zero x^k and x^0 terms
 - All 2-bit errors if $C(x)$ has a factor with at least three terms
 - Any odd number of errors if $C(x)$ has $(x + 1)$ as a factor
 - Any burst error $< k$ bits
- There are standardized polynomials of different degree that are known to catch many errors
 - Ethernet CRC-32:
100000100110000010001110110110111

Error Correction

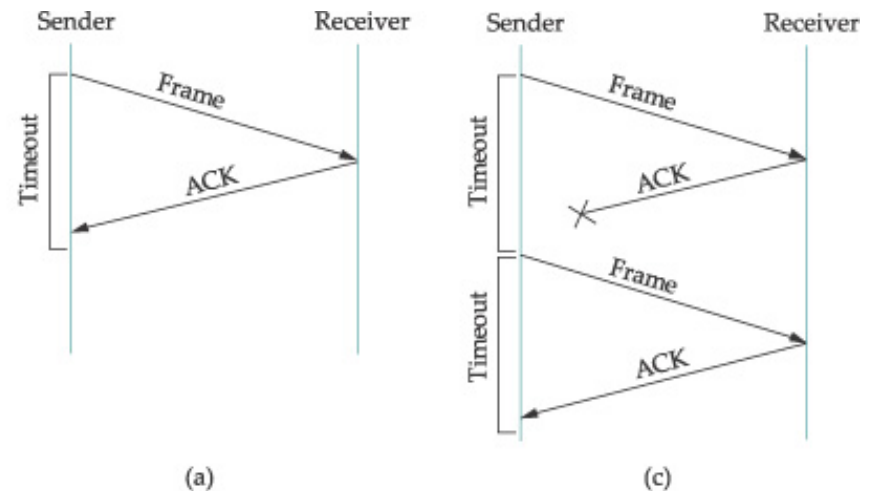
- Two strategies to correct errors:
 - Detect and retransmit, or Automatic Repeat reQuest. (ARQ)
 - Error correcting codes, or Forward Error Correction (FEC)
- Retransmissions typically at higher levels (Network+). Why?
- Question: Which should we choose?

Retransmissions vs. FEC

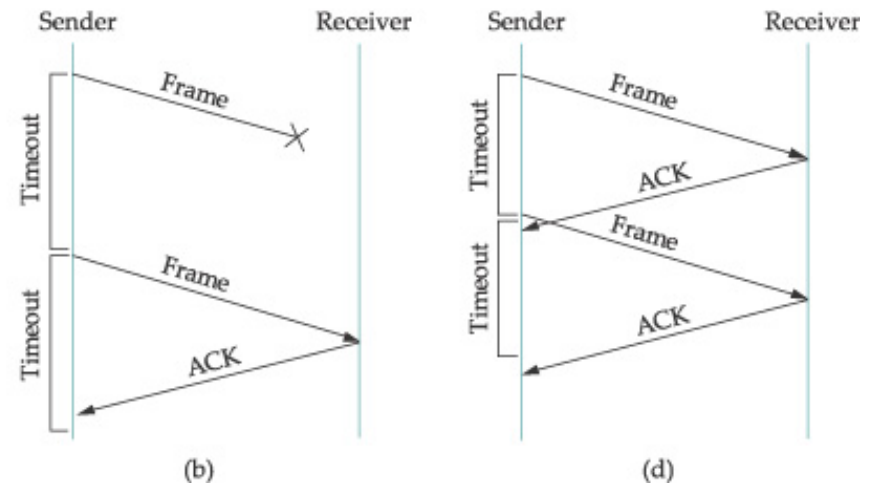
- The better option depends on the kind of errors and the cost of recovery
- Example: Message with 1000 bits, Prob(bit error) 0.001
 - Case 1: random errors
 - Case 2: bursts of 1000 errors
 - Case 3: real-time application (teleconference)

ARQ: Stop-and-Wait

- Idea: transmit and wait for an ACK
 - Sender sets a timer
 - He retransmits if doesn't hear an ACK before the timer expires
 - What should the timeout be?
 - Do we need sequence numbers?





- Efficiency?
 - E.g., 1.5 Mbps link with 45ms RTT, 1500 byte frame → $1500 * 8 / .045 = 267$ Kbps



ARQ: Sliding Window

- Idea: have more than one outstanding packet
- $LFS - LAR \leq SWS$, $LAF - LFR \leq RWS$
- How big should SWS and RWS be?
- Sequence numbers?
- Flow Control?
- Cumulative ACKS
- Selective ACKs? NACKs?



FEC: Reed-Solomon / BCH Codes

- Developed to protect data on magnetic disks
- Used for CDs and cable modems too
- Property: $2t$ redundant bits can correct $\leq t$ errors
- Mathematics somewhat more involved ...

Key Concepts

- Senders “frame” messages with sentinels, length fields, and clock synch, so receivers can determine where they start and end
- Redundant bits are added to messages to protect against transmission errors.
- Two recovery strategies are retransmissions (ARQ) and error correcting codes (FEC)