


CSE 561 – Lecture 4

Reliability

David Wetherall
djw@cs.washington.edu

Automatic Repeat Request (ARQ)




- Packets can be corrupted or lost. How do we add reliability?
- Acknowledgments (ACKs) and retransmissions after a timeout
- ARQ is generic name for protocols based on this strategy

djw // CSE/EE 461, Winter 2003

L13.2

The Need for Sequence Numbers




- In the case of ACK loss (or poor choice of timeout) the receiver can't distinguish this message from the next
 - Need to understand how many packets can be outstanding and number the packets; here, a single bit will do

djw // CSE/EE 461, Winter 2003

L13.3

Stop-and-Wait

- Only one outstanding packet at a time
- Also called alternating bit protocol



djw // CSE/EE 461, Winter 2003

L13.4

Limitation of Stop-and-Wait




- Lousy performance if wire time \ll prop. delay
 - How bad? You do the math
- Want to utilize all available bandwidth
 - Need to keep more data "in flight"
 - How much? Remember the bandwidth-delay product?
- Leads to Sliding Window Protocol

djw // CSE/EE 461, Winter 2003

L13.5

Sliding Window – Sender




- Window bounds outstanding data
 - Implies need for buffering at sender
- "Last" ACK applies to in-order data
- Sender maintains timers too
 - Go-Back-N: one timer, send all unacknowledged on timeout
 - Selective Repeat: timer per packet, resend as needed

djw // CSE/EE 461, Winter 2003

L13.6


Sliding Window – Timeline



djw // CSE/EE 461, Winter 2003

L13.7

Sliding Window – Receiver



- Receiver buffers too:
 - data may arrive out-of-order
 - or faster than can be consumed (flow control)
- Receiver ACK choices:
 - Individual, Cumulative (TCP), Selective (newer TCP), Negative

djw // CSE/EE 461, Winter 2003

L13.8

Connection Establishment


- Both sender and receiver must be ready before we start to transfer the data
 - Sender and receiver need to agree on a set of parameters
 - e.g., the Maximum Segment Size (MSS)
- This is signaling
 - It sets up state at the endpoints
 - Compare to “dialing” in the telephone network
- In TCP a Three-Way Handshake is used

djw // CSE/EE 461, Winter 2003

L13.9

Three-Way Handshake

- Opens both directions for transfer



djw // CSE/EE 461, Winter 2003

L13.10

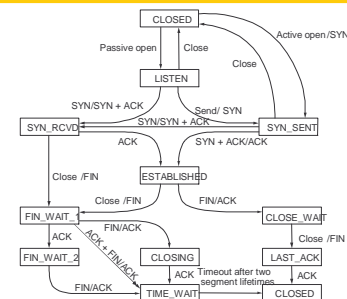
Some Comments

- We could abbreviate this setup, but it was chosen to be robust, especially against delayed duplicates
 - Three-way handshake from Tomlinson 1975
- Choice of changing initial sequence numbers (ISNs) minimizes the chance of hosts that crash getting confused by a previous incarnation of a connection
- But with random ISN it actually proves that two hosts can communicate
 - Weak form of authentication

djw // CSE/EE 461, Winter 2003

L13.11


TCP State Transitions



djw // CSE/EE 461, Winter 2003

L13.12

Again, with States



djw // CSE/EE 461, Winter 2003

L13.13


Connection Teardown

- Orderly release by sender and receiver when done
 - Delivers all pending data and "hangs up"
- Cleans up state in sender and receiver
- TCP provides a "symmetric" close
 - both sides shutdown independently

djw // CSE/EE 461, Winter 2003

L13.14

TCP Connection Teardown



djw // CSE/EE 461, Winter 2003

L13.15

The TIME_WAIT State


- We wait 2MSL (two times the maximum segment lifetime of 60 seconds) before completing the close
- Why?
 - ACK might have been lost and so FIN will be resent
 - Could interfere with a subsequent connection

djw // CSE/EE 461, Winter 2003

L13.16

UDP Checksum

- UDP includes optional protection against errors
 - Checksum intended as an end-to-end check on delivery
 - So it covers data, UDP header, and IP pseudoheader



djw // CSE/EE 461, Winter 2003

L13.17

Errors and Redundancy

- Noise can flip some of the bits we receive
 - We must be able to detect when this occurs!
- Basic approach: add redundant data
 - Error detection codes allow errors to be recognized
 - Error correction codes allow errors to be repaired too

djw // CSE/EE 461, Winter 2003

L13.18

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

djw // CSE/EE 461, Winter 2003

L13.19

Detection vs. Correction

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

djw // CSE/EE 461, Winter 2003

L13.20

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)

djw // CSE/EE 461, Winter 2003

L13.21

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 à 000

djw // CSE/EE 461, Winter 2003

L13.22

Parity

- Start with n bits and add another so that the total number of 1s is even (even parity)
 - e.g. 0110010 à 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

djw // CSE/EE 461, Winter 2003

L13.23

2D Parity

- Add parity row/column to array of bits

- Detects all 1, 2, 3 bit errors, and many errors with >3 bits.
- Corrects all 1 bit errors

```
↓
0101001 1
1101001 0
1011110 1
0001110 1
0110100 1
1011111 0
→ 1111011 0 ←
↑
```

djw // CSE/EE 461, Winter 2003

L13.24

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

djw // CSE/EE 461, Winter 2003

L13.25

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 $x^7 + x^4 + x^3 + x^1$

djw // CSE/EE 461, Winter 2003

L13.26

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

djw // CSE/EE 461, Winter 2003

L13.27

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 ...

djw // CSE/EE 461, Winter 2003

L13.28