

What you saw in DigiComms I

- What it looks like on the wire
 - Systematic or random noise
 - Attenuation
 - Signal to noise ratio
- Error Detection and Correction
 - Forward Error Correction (FEC)
 - Parity
 - Block codes
 - CRCs

What you saw in DigiComms I (2)

Protocols

- Retransmit packets on NACKs or timeouts
- Why single Automatic Repeat Request (ARQ) is poor
 - Want to maximise traffic in transit (within the network)
- Continuous ARQ: using a window
 - Go back N
 - Selective retransmission

CRC

Detects

- All single bit errors
- Almost all 2-bit errors
- Any odd number of errors
- All bursts up to M, where generator length is M
 - + Networks errors: usually burst, not randomly distributed
- Longer bursts with probability 2^{-m}

2



Software schemes

Efficiency is important

- Touch each data byte only once
- CRC
- TCP/UDP/IP
 - All use same scheme
 - Treat data bytes as 16-bit integers
 - Add with end-around carry
 - One's complement = checksum
 - Catches all 1-bit errors
 - Longer errors with probability 1/65536

Packet errors

- Different from bit errors
 - Types
 - + Not just erasure, but also duplication, insertion, etc.
 - Correction
 - + Retransmission, instead of redundancy

Types of packet errors

Loss

- Due to uncorrectable bit errors
- Buffer loss on overflow
 - + Especially with bursty traffic
 - For the same load, the greater the burstiness, the more the loss
 - + Loss rate depends on burstiness, load, and buffer size
- Fragmented packets can lead to error multiplication
 - + Longer the packet, more the loss

6



Packet error detection and correction

- Detection
 - Sequence numbers
- Timeouts
- Correction
 - Retransmission

Sequence numbers

- In each header
- Incremented for non-retransmitted packets
- Sequence space
 - Set of all possible sequence numbers (seq #s)
 - For a 3-bit sequence number, space is {0,1,2,3,4,5,6,7}

Using sequence numbers

- Loss
 - Gap in sequence space allows receiver to detect loss
 - + e.g. received 0,1,2,5,6,7 => lost 3,4
 - ACKs carry cumulative sequence number
 - Redundant information
 - If no ACK for a while, sender suspects loss
- Reordering
- Duplication
- Insertion

11

- If the received sequence number is "very different" from what is expected
 - + More on this later...

10



In general 2^{sequence_size > R(2 MPL + T + A)}

Maximum Packet Lifetime (MPL)

- How can we bound it?
- Generation time in header
 - Too complex!
- Counter in header decremented per hop
 - Crufty, but works
 - Used in the Internet
 - Time To Live (TTL)
 - Assumes a maximum network diameter, and a limit on forwarding time



Packet insertion

- Receiver should be able to distinguish packets from other connections
- Why?

13

- Receive packets on VCI 1
- Connection closes
- New connection also with VCI 1
- Delayed packet arrives
- Could be accepted
- Solution
 - Flush packets on connection close
 - · Can't do this for connectionless networks like the Internet

14



- TCP/UDP packets carry source IP, destination IP, source port number, destination port number
- How we can have insertion?
 - Host A opens connection to B, source port 123, destination port 456
 - Transport layer connection terminates
 - New connection opens, A and B assign the same port numbers

17

19

- Delayed packet from old connection arrives
- Insertion!

Solutions

- Per-connection incarnation number
 - Incremented for each connection from each host
 - Takes up header space
 - + On a crash, we may repeat
 - Need stable storage, which is expensive
- Reassign port numbers only after 1 MPL
 - Needs stable storage to survive crash

Solutions (cont.)

- Assign port numbers serially: new connections have new ports
 - Unix starts at 1024
 - This fails if we wrap around within 1 MPL
 - Also fails of computer crashes and we restart with 1024
- Assign initial sequence numbers serially
 - New connections may have same port, but sequence # differs
 - Fails on a crash
- Wait 1 MPL after boot up (30s to 2 min)
 - This flushes old packets from network
 - Used in most Unix systems

3-way handshake

- Standard solution, then, is
 - Choose port numbers serially
 - Choose initial sequence numbers from a clock
 - Wait 1 MPL after a crash
- Needs communicating ends to tell each other initial sequence number
- Easiest way is to tell this in a SYNchronize packet (TCP) that starts a connection
- 2-way handshake

18



- Problem really is that SYNs themselves are not protected with sequence numbers
- 3-way handshake protects against delayed SYNs
- (Server passive opens a port e.g. 80)
- Client active opens connection with SYN packet
 Client's sequence number is 'c'
- Server replies with SYN-ACK ACK=(c+1), SYN=(s)
 Server's sequence number is 's'
- Client sends ACK to server ACK=(s+1).
- Recall that TCP sequence numbers are per data byte

Loss detection

- At receiver, from a gap in sequence space
 Send a NACK to the sender
- At sender, by looking at cumulative ACKs, and timing out if no ACK for a while
 - Need to choose timeout interval

NACKs

- Sounds good, but does not work well
 - Adds extra load during times of loss:
 - Probably not helpful if loss is from congestion!
 - At least NACKs travel in the reverse direction...
- If NACK is lost, receiver must retransmit it
 - Moves timeout problem to receiver
- So we need timeouts anyway

Timeouts

21

23

- Set timer on sending a packet
- If timer goes off, and no ACK, resend
- How to choose timeout value?
- Intuition: expect a reply in about one round trip time (RTT)
- Aside: Timeouts usually preclude freeing their related resources
 - What about faked requests?
- SYN attack
 - Forge an unused IP address
 - Send a long sequence of SYN packets to server
 - Up to four minutes' resource allocation per SYN while server retries its SYN-ACKs
 - Without protection, server falls over

22







Error control window

Set of packets sent, but not ACKed
1 2 3 4 5 6 7 8 9 (original window)
1 2 3 4 5 6 7 8 9 (receive ACK for 3)
1 2 3 4 5 6 7 8 9 (send 8)
May want to restrict max size = window size
Sender blocked until ACK comes back



Selective retransmission

- Somehow find out which packets lost, then only retransmit them
- How to find lost packets?
 - Each ACK has a bitmap of received packets
 - + e.g. cum_ACK = 5, bitmap = 101 \rightarrow received 5 and 7, but not 6
 - + Wastes header space
 - Sender periodically asks receiver for bitmap
 - Fast retransmit

31

30

Fast retransmit

- Assume cumulative ACKs
- If sender sees repeated cumulative ACKs, packet likely lost
- **1**, 2, 3, 4, 5, 6
- **1**, 2, 3, 3, 3
- Send cumulative_ACK + 1 = 4
- Used in TCP

SMART

- ACK carries cumulative sequence number
- Also sequence number of packet causing ACK
 - 1 2 3 4 5 6 7 packets sent
 - 1 2 3 3 3 3 received ACK (cumulative)
 - 1 2 3 5 6 7 received ACK (for packet)
- Sender creates bitmap
- No need for timers!

33

If retransmitted packet lost, periodically check if cumulative ACK increased.