

Internet Technology

05. Transport Layer

Paul Krzyzanowski
Rutgers University
Spring 2016

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 1

Transport Layer

- Transport Layer
 - Provides logical communication channels between apps
- Transport layer managed by end systems
 - Routers are unaware; they provide network layer services
- Multiple transport protocols available
 - Under IP: TCP, UDP, SCTP, and more

Internet Protocol Layers

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 2

Transport Layer

- Network Layer
 - Logical connection between hosts
- Transport Layer
 - Logical connection between processes
 - Transport layer **multiplexing & demultiplexing**
- Most common transport-layer protocols in IP: TCP & UDP
 - UDP: unreliable data transfer
 - TCP
 - Reliable data transfer
 - In-order delivery
 - Flow control
 - Congestion control

Internet Protocol Layers

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 3

Today, we'll discuss

- Transport layer multiplexing/demultiplexing
- Reliable data transfer

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 4

Transport Layer Multiplexing & Demultiplexing

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 5

Transport Layer Multiplexing

- Problem:
 - Multiple communication channels over one network link
 - This is a problem whenever a protocol at one layer is used by multiple protocols or communication sessions at a higher layer

Logical view of four transport layer communication streams

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 6

Transport Layer Multiplexing

- Problem:
 - Multiple communication channels over one network link
 - This is a problem whenever a protocol at one level is used by multiple protocols or or communication session at one
- Need to identify which segment belongs to which channel

Logical view at the network layer

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 7

Multiplexing & Demultiplexing

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 8

How is it done?

- Transport layer protocols in IP have **port numbers**
 - 16 bit integers (0 .. 65535)
 - IP header (network layer) has **source address, destination address**
 - TCP/UDP headers (transport layer) have **source port, destination port**
- Each socket is uniquely identified in the operating system
- Before a socket can be used, it is created & named
 - **socket** system call creates a unique socket
 - **bind** system call associates a local address with the socket
 - With an address of INADDR_ANY, the socket is associated with ALL local interfaces
 - With a port of 0, the OS assigns a random unused port number to the socket

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 9

UDP multiplexing & demultiplexing

- A UDP socket is identified by its port number
- All UDP segments addressed to a specific port # will be delivered to the socket identified by that port number
 - A socket will request data via `recv()`, `recvfrom()`, or `recvmsg()` system calls
 - OS looks for a UDP socket with a matching destination port: hash table of socket structures; hash key created from UDP destination port
- Limited demultiplexing
 - Segments addressed to the same (*host, port*) from different processes or different systems will be delivered to the same socket!
 - The receiver can get the source address & port to know how to address reply messages

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 10

Why use UDP?

- Control the timing of data
 - A UDP segment is passed to the network layer immediately for transmission
 - TCP uses congestion control to delay transmission
- Preserve message boundaries
 - With TCP, multiple small messages may be consolidated into one TCP segment
- No connection setup
 - TCP requires a three-way handshake to establish a connection
- No state to keep track of
 - Less memory, easier fault recovery, simple load balancing
- Less network overhead
 - 8-byte header instead of TCP's 20-byte header

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 11

UDP Structure

- Defined in RFC 768
- Eight byte header

32 bits 4 bytes	
Source Port #	Dest. Port #
Length	Checksum
Application Data	

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 12

UDP Structure in context

Eight byte header within a 20 byte IP header

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 13

UDP Checksum

- IP does not guarantee error-free packet delivery
- The UDP header contains a 16-bit checksum
 - Checks for data corruption
- Checksum is generated by the sender and validated only by the receiver only. segments with bad checksums are simply dropped

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 14

UDP Checksum Calculation

- Sender**
 - Iterate over 16-bit words in the Pseudo header + UDP segment
 - UDP checksum field = 0
 - Create a **one's complement checksum**
 - Add two 16-bit values. If overflow, add 1 to the result
 - Do this for all the data you need to checksum
 - Invert the bits of the result to get the checksum value
- Receiver**
 - Perform the same one's complement sum on all data *including* the checksum field
 - The result should be all 1s (0xffff)

The same checksum calculation is used for the IP header, UDP header, & TCP header

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 15

One's Complement Checksum Example

- How to compute a One's complement
 - Sum the numbers
 - Add any overflow carry to the result
- Create checksum for:


```

0110 1011 0000 1010
1011 0001 1100 1000
1100 0000 1111 0101
1100 0000 1111 0101
            
```

```

0110 1011 0000 1010
+1011 0001 1100 1000
=10001 1100 1101 0010
+
=0001 1100 1101 0011
+1100 0000 1111 0101
=1101 1101 1100 1000
            
```
- Then invert the bits


```

~1101 1101 1100 1000
=0010 0010 0011 0111 ← checksum
            
```

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 16

One's Complement Checksum Example

- Validate**
 - Sum the numbers, including the checksum

```

0110 1011 0000 1010
1011 0001 1100 1000
1100 0000 1111 0101
add the checksum → 0010 0010 0011 0111
            
```

```

0110 1011 0000 1010
+1011 0001 1100 1000
=10001 1100 1101 0010
+
=0001 1100 1101 0011
+1100 0000 1111 0101
=1101 1101 1100 1000
+0010 0010 0011 0111
=1111 1111 1111 1111 ← add checksum
            
```

- A result of all 1's (= -0) means the transmission was good

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 17

TCP multiplexing & demultiplexing

- Every TCP socket is identified by:

(source address, destination address, source port, destination port)
- A TCP socket has a state:
 - LISTEN**: the socket is used only for accepting connections
 - ESTABLISHED**: the socket is connected
 - Other states that we'll ignore for now:
 - Connection setup:
 - SYN_SENT: trying to establish a connection
 - SYN_RCVD: received a connection request
 - Connection teardown:
 - FIN_WAIT_1: socket has been closed by the local application, no acknowledgement from remote
 - FIN_WAIT_2: socket has been closed by the local application, remote acknowledged the closing
 - CLOSING: socket has been closed by the local & remote apps; remote has not acknowledged close
 - TIME_WAIT: connections closed; waiting to be sure that the remote side received the last ACK
- Let's look at an example

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 18

Server: Create a new socket

```
svr = socket(AF_INET, SOCK_STREAM, 0);
```

Address family: Internet (IPv4)
Type: 'Stream' - connection-oriented (TCP)

Create a new socket at the server: it has no addresses so far

Client (135.10.10.1)				
LocalAddr	LocalPort	Remote Addr	RemPort	State

Server (192.115.8)				
LocalAddr	LocalPort	Remote Addr	RemPort	State

SVR →

N.B.: We refer to a socket table here for convenience but it is just a logical construct. The actual implementation is operating-system specific but this data is generally stored in a list of socket buffer structures. On Linux, for example, the kernel function tcp_v4_lookup will search for either a listening or an established socket with specific addresses and ports (see net/ipv4/tcp_ipv4.c, around line 507)

March 29, 2016

CS 352 © 2013-2016 Paul Krzyzanowski

19

Server: Bind - assign a local address

```
bind(svr);
```

Assign a local address (INADDR_ANY) and port (1234) to the socket

Client (135.10.10.1)				
LocalAddr	LocalPort	Remote Addr	RemPort	State

Server (192.115.8)				
LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			

SVR →

March 29, 2016

CS 352 © 2013-2016 Paul Krzyzanowski

20

Server: Make it a listening socket

```
listen(svr, 10);
```

Set the state of the socket to listen. This socket can only be used to accept connections

Client (135.10.10.1)				
LocalAddr	LocalPort	Remote Addr	RemPort	State

Server (192.115.8)				
LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN

SVR →

March 29, 2016

CS 352 © 2013-2016 Paul Krzyzanowski

21

Server: Wait for a connection

```
snw = accept(svr);
```

Wait for an incoming connection on this socket

Client (135.10.10.1)				
LocalAddr	LocalPort	Remote Addr	RemPort	State

Server (192.115.8)				
LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN

SVR →

March 29, 2016

CS 352 © 2013-2016 Paul Krzyzanowski

22

Client: Create a new socket

```
s = socket();
```

Create a new socket at the client: no addresses so far

Client (135.10.10.1)				
LocalAddr	LocalPort	Remote Addr	RemPort	State

← s

Server (192.115.8)				
LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN

SVR →

March 29, 2016

CS 352 © 2013-2016 Paul Krzyzanowski

23

Client: Assign a local address & port

```
bind(s);
```

Assign any local address (INADDR_ANY) and have the OS pick a port (port=0)

Client (135.10.10.1)				
LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801			

← s

Server (192.115.8)				
LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN

SVR →

March 29, 2016

CS 352 © 2013-2016 Paul Krzyzanowski

24

Client: Connect to the server

`connect(s, dest_addr);`

Connect to address 192.11.5.8, port 1234

Client (135.10.10.1)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801			

← s

connection

Server (192.115.8)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN
192.11.5.8	1234	135.10.10.1	7801	SYN_RCVD

svr →
snew →

- Send a connection establishment request to address 192.11.5.8, port 1234 (TCP segment to port 1234 with a connection setup bit set, we'll look at the exact handshake later)
- On the server, search the table for a LISTEN socket where packets destination addr == table's local addr (0.0.0.0 matches any incoming addr) packets destination port == table's local port
- Create a new socket for the connection

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 25

Client: Complete the connection

`connect(s, dest_addr);`

Server acknowledges the connection; Client fills in the entry

Client (135.10.10.1)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801	192.11.5.8	1234	ESTABLISHED

← s

Server (192.115.8)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN
192.11.5.8	1234	135.10.10.1	7801	ESTABLISHED

svr →
snew →

Now we can talk!

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 26

Communicate

Client-to-server communication

Server finds socket by searching for a TCP socket with these properties:

- Status == ESTABLISHED
- IP src addr == remote addr
- TCP src port == remote port
- IP destaddr == local addr
- TCP destport == local port

Client (135.10.10.1)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801	192.11.5.8	1234	ESTABLISHED

← s

Server (192.115.8)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN
192.11.5.8	1234	135.10.10.1	7801	ESTABLISHED

svr →
snew →

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 27

Communicate

Server-to-client communication

Client finds socket by searching for a TCP socket with these properties:

- Status == ESTABLISHED
- IP src addr == remote addr
- TCP src port == remote port
- IP destaddr == local addr
- TCP destport == local port

Client (135.10.10.1)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801	192.11.5.8	1234	ESTABLISHED

← s

Server (192.115.8)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN
192.11.5.8	1234	135.10.10.1	7801	ESTABLISHED

svr →
snew →

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 28

Two clients sharing the same port

Different source address disambiguates the sessions

Client (135.10.10.2)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801	192.11.5.8	1234	ESTABLISHED

Client (135.10.10.1)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801	192.11.5.8	1234	ESTABLISHED

Server (192.115.8)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN
192.11.5.8	1234	135.10.10.1	7801	ESTABLISHED
192.11.5.8	1234	135.10.10.1	7801	ESTABLISHED

svr →
snew1 →
snew2 →

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 29

Two endpoints sharing the same address

The OS will not allow two sockets to share the same port on one client

Client (135.10.10.1)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	7801	192.11.5.8	1234	ESTABLISHED
0.0.0.0	7802	192.11.5.8	1234	ESTABLISHED

Server (192.115.8)

LocalAddr	LocalPort	Remote Addr	RemPort	State
0.0.0.0	1234			LISTEN
192.11.5.8	1234	135.10.10.1	7801	ESTABLISHED
192.11.5.8	1234	135.10.10.1	7802	ESTABLISHED

svr →
snew1 →
snew2 →

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 30

Reliable Data Transfer

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 31

Reliable Data Transfer (RDT) Goal

Develop a protocol for transmitting data reliably over an unreliable network

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 32

RDT over a reliable channel

- Assume the channel is reliable
- Trivial – nothing to do!

Here's the finite state machine (FSM):

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 33

RDT over a channel with bit errors

- All packets are received
- Some might be corrupt
- Approach
 - Acknowledge each packet
 - Positive acknowledgement (ACK): "I got it; looks good!"
 - Negative acknowledgement (NAK): "Please repeat"
 - Sender retransmits a packet if it receives a NAK
- ARQ (Automatic Repeat reQuest)
 - Set of protocols that use acknowledgements & retransmission

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 34

We need to support three capabilities

- Error detection**
 - How do we know if the packet is corrupt?
 - Use a checksum (**error detecting code**)
- Receiver feedback**
 - The receiver will acknowledge each packet with an ACK or NAK
- Retransmission**
 - If a sender gets a NAK, the packet will be retransmitted

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 35

RDT over a channel with bit errors

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 36

Stop-and-wait

- The sender cannot send any data until it receives an ACK for the previously sent packet
- This type of protocol is a **stop-and-wait** protocol

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 37

What about a corrupted ACK/NAK message?

- The sender does not know whether the last packet was received correctly or not
- We can
 - Have the sender send a "please repeat" in response to a corrupt ACK/NAK
 - But what if that gets corrupted?
 - Add a robust error correcting code
 - Works for a channel that does not lose data
 - Resend the data in response to a corrupted ACK/NAK
 - Duplicate packets may be received
 - Receiver needs to distinguish between new data & a retransmission
 - Use a **sequence number**. Here, we only need a 1-bit number.

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 38

A 1-bit sequence number

Sequence bit flip-flops between consecutive messages

Alternating bit protocol

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 39

RDT over a channel with bit errors

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 40

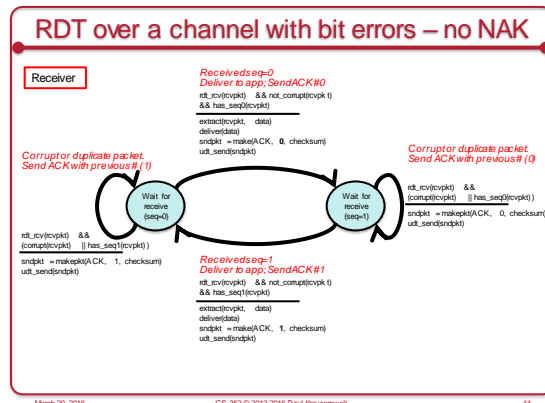
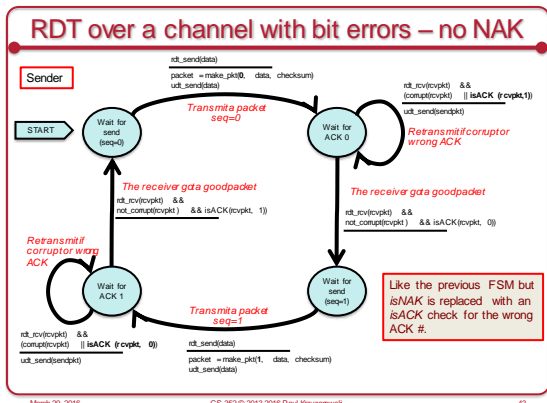
RDT over a channel with bit errors

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 41

RDT over a channel with bit errors

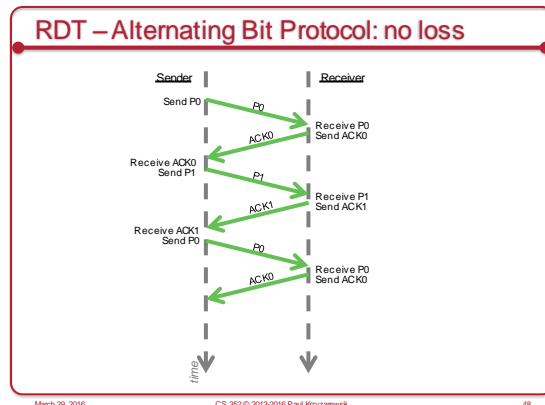
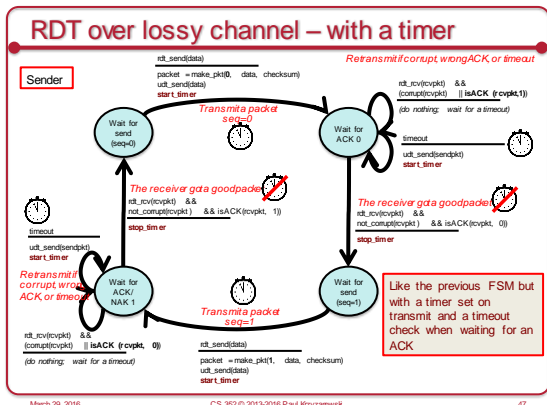
- If a corrupted packet is received
 - Send a NAK
- If a duplicate packet is received
 - Send an ACK since we already processed the packet
- We can get rid of NAKs
 - Send an ACK for the last correctly received packet
 - If a sender receives duplicate ACKs, it knows that the previous packet has not been received correctly
 - Modify protocol: add **sequence numbers** to ACKs

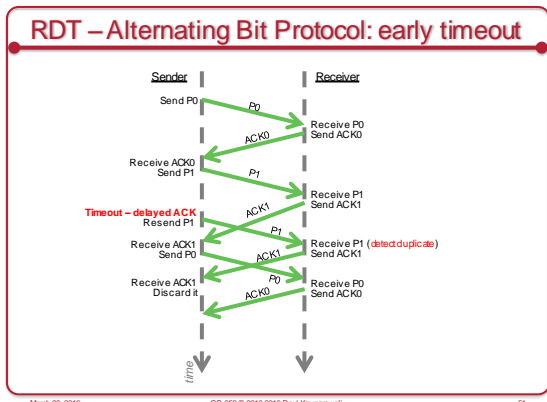
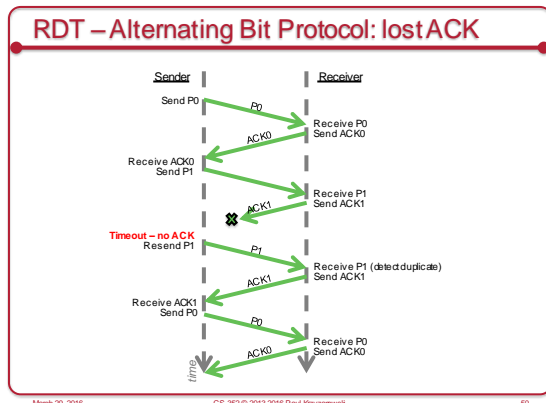
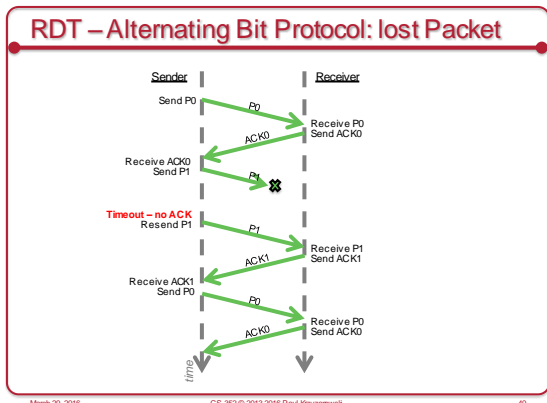
March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 42



- ### RDT over a lossy channel
- We considered only bit errors
 - Packets were always delivered
 - How do we detect & deal with packet loss?

- ### Dealing with packet loss
- Burden of detection & recovery is on sender
 - If sender's packet is lost OR receiver's ACK is lost
 - Sender will not get a reply from the receiver
 - Approach
 - Introduce a **countdown timer**. Set the timer at transmit
 - If time-out and no reply retransmit
 - How long to wait? Maximum round-trip delay?
 - Long wait until we initiate error recovery
 - Pick a "likely loss" time
 - Retransmit if no response within that time
 - Introduces possibility of **duplicate packets**
 - But we already know how to deal with them





Network utilization with stop-and-wait

- A **stop-and-wait** protocol gives us **horrible network utilization**
- Consider
 - Cross-country link \Rightarrow Round-trip propagation delay (RTT) = 30 ms
 - Assume 1 Gbps link (ignore router delays), $R = 10^9$ bits/second
 - Assume 1,000-byte packets ($L = 8,000$ bits)
 - Time to transmit the packet: $d_{trans} = L / R = 8,000 / 10^9 = 8 \mu s$
- With a stop-and-wait protocol
 - one-way delay = $d_{trans} + d_{prop} = (30 \text{ ms} + 2) + 8 \mu s = 15.008 \text{ ms}$
 - Assume ACK packets are tiny; one-way delay for ACK packet = 15 ms
 - ACK is received at $15.008 + 15 = 30.008 \text{ ms}$
 - Next packet can be sent $(15.008 + 15) = 30.008 \text{ ms}$ after the first one
 - Utilization = fraction of time sender is sending bits into the channel

$$U = \frac{L/R}{RTT + (L/R)} = \frac{0.008}{30.008} = 0.00027 = 0.027\%$$

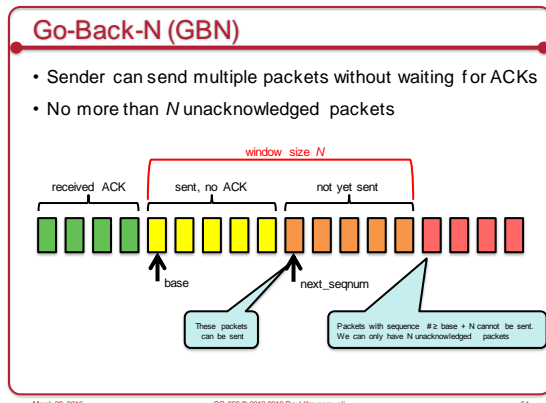
The sender can transmit 1,000 bytes in 30,008 ms: 267 kbps on a 1 Gbps link!

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 52

Improve Network Utilization: Pipelining

- Don't wait for an acknowledgement before sending the next packet
- But then we need to
 - Increase the range of sequence numbers
 - Each in-transit packet needs a unique number
 - Hold on to unacknowledged packets at sender
 - Hold on to out-of-sequence packets at receiver
- Two approaches for pipelined error recovery
 - Go-Back-N**
 - Selective Repeat**

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 53



Go-Back-N (GBN)

- Sender can send multiple packets without waiting for ACKs
- No more than N unacknowledged packets

GBN = Sliding Window Protocol

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 55

Go-Back-N (GBN)

- Sender can send multiple packets without waiting for ACKs
- No more than N unacknowledged packets

GBN = Sliding Window Protocol

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 56

Go-Back-N (GBN)

- Sender can send multiple packets without waiting for ACKs
- No more than N unacknowledged packets

GBN = Sliding Window Protocol

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 57

Go-Back-N (GBN)

- Sender can send multiple packets without waiting for ACKs
- No more than N unacknowledged packets

GBN = Sliding Window Protocol

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 58

Sequence numbers

A sequence number will take up a fixed #, k , of bits in the header

- Range of sequence numbers is $0 \dots 2^k - 1$
- Modulo 2^k arithmetic: $2^k - 1$ increments to 0

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 59

Extended FSM for a GBN sender

```

Sender
    Send data if it's in the window (we can have at most N unacknowledged packets)
    ut_send(data)
    if (next_seqnum < base+N) { // there's room in the window
        sendpkt(next_seqnum) // make_pkt(next_seqnum, data, checksum)
        ut_send(sendpkt) end a seqnum
        if (base == next_seqnum)
            start_timer
            next_seqnum++
    } else {
        // else, do nothing if cannot send
        // else, do nothing if cannot send
    }

    ignore corrupted ACKs
    ut_rcv(evpkt) && comp(evpkt)

    Wait
    start_timer
    for i = base; i < next_seqnum; i++
        ut_send(sendpkt[i])

    ut_rcv(evpkt) && not_comp(evpkt)
    base = get_ackno(evpkt)+1
    if (base == next_seqnum)
        stop_timer // we have the latest ACK
    else
        start_timer // still waiting for ACKs
    
```

Go Back-N: Timeout means resend all unacknowledged packets

Calculation acknowledgment: Receipt of a sequence number n ACK means that all packets up to and including n have been received

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 60

Extended FSM for a GBN receiver

Receiver

```

We received a good packet with the expected sequencenumber
rtt_send(rcvpkt) && !not_computed(rcvpkt) && has_seqnum(rcvpkt, expected_seqnum)
extract(rcvpkt, data)
deliver(data) // give it to the app
sndpkt = makepkt(expected_seqnum, ACK, checksum)
utl_send(sndpkt) // send the ACK to the sender
expected_seqnum++
    
```

Initialize
 expected_seqnum = 1
 sndpkt = makepkt(0, ACK, checksum)

The receiver discards out-of-order packets

If packet n is lost and $n+1$ arrives, the receiver does not buffer packet $n+1$. The sender will retransmit all unacknowledged packets (go back N).

The receiver has to only keep track of the next sequence number.

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 61

Selective Repeat

- Problem with Go-Back-N**
 - With a large window size and large delays, many packets can be in the pipeline
 - A single error can cause GBN to retransmit many packets (all that are unacknowledged)
 - If $P(\text{channel error})$ increases, the pipeline can become filled with excess retransmissions
- Selective Repeat Protocol**
 - Retransmit only those packets that were lost or corrupted
 - Receiver must acknowledge each correctly received packet
 - Even if it is out of order
 - Out of order packets must be buffered
 - Window size $N =$ limit of number of outstanding packets
 - But some packets in the window may be acknowledged
 - The window slides when the earliest packet in the window is acknowledged

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 62

Selective Repeat Windows

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 63

Selective Repeat: sender operation

- Send requests from application**
 - Check next available sequence #
 - If no room in window, reject (or buffer)
 - Else send the packet (with sequence #)
- Timeout**
 - Each packet has its own timer
 - Retransmit only the specific packet on timeout
- ACK received**
 - If packet is within window
 - Mark packet as received
 - If $\text{sequence \#} == \text{send_base}$ advance the base (start of window) to the next unacknowledged sequence number

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 64

Out-of-order ACKs

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 65

Out-of-order ACKs

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 66

Selective Repeat: receiver operation

- Good packet with seq # in $[rcv_base, rcv_base+N-1]$
 - Packet is within the receiver's window
 - Send ACK for that sequence #
 - If sequence # == rcv_base
 - Deliver packet to app and deliver all successive packets that have been received
 - Adjust start of window (rcv_base)
- Good packet with seq # in $[rcv_base-N, rcv_base-1]$
 - Packet is within the *before* receiver's window
 - We already saw it – but send ACK anyway
- Anything else
 - Ignore the packet

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 67

Selective Repeat: receiving packets

Receiver's view

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 68

Selective Repeat: receiving packets

Receiver's view

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 69

The end

March 29, 2016 CS 352 © 2013-2016 Paul Krzyzanowski 70