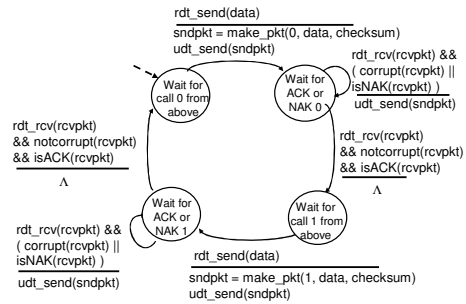


Reliable Data Transfer

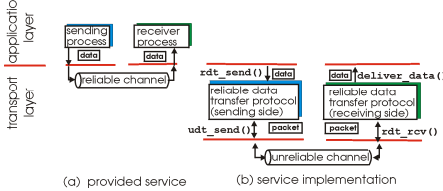
EECS 122
Feb. 9, 2006

rdt2.1: sender, handles garbled ACK/NAKs



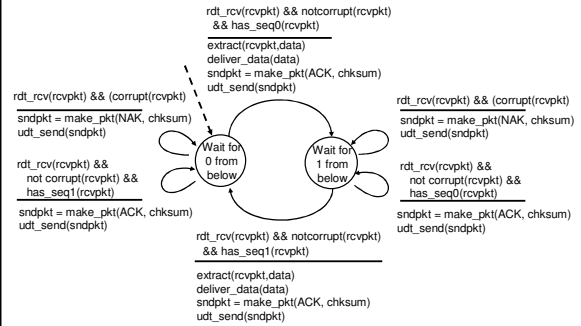
Recap: Reliable data transfer

- Goal: provide reliable packet delivery over unreliable channels.



- We have considered bit flips as the only channel impairment
- Proposed a stop-and-wait protocol based on ACK/NACK

rdt2.1: receiver, handles garbled ACK/NAKs



rdt2.0 has a fatal flaw!

What happens if ACK/NAK corrupted?

- sender doesn't know what happened at receiver!
- can't just retransmit: possible duplicate

Handling duplicates:

- sender retransmits current pkt if ACK/NAK garbled
- sender adds *sequence number* to each pkt
- receiver discards (doesn't deliver up) duplicate pkt

rdt2.1: discussion

Sender:

- seq # added to pkt
- two seq. #'s (0,1) will suffice. Why?
- must check if received ACK/NAK corrupted
- twice as many states
 - state must "remember" whether "current" pkt has 0 or 1 seq. #

Receiver:

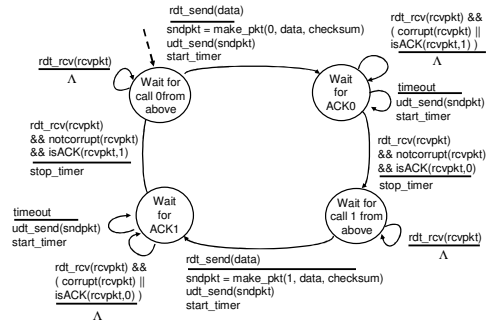
- must check if received packet is duplicate
 - state indicates whether 0 or 1 is expected pkt seq #
- note: receiver can *not* know if its last ACK/NAK received OK at sender

rdt2.2: a NAK-free protocol

- same functionality as rdt2.1, using ACKs only
- instead of NAK, receiver sends ACK for last pkt received OK
 - receiver must *explicitly* include seq # of pkt being ACKed
- duplicate ACK at sender results in same action as NAK: *retransmit current pkt*

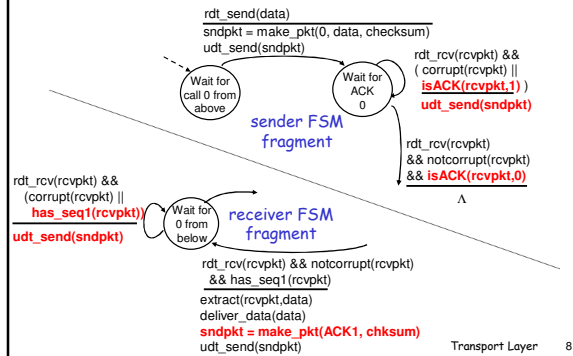
Transport Layer 7

rdt3.0 sender



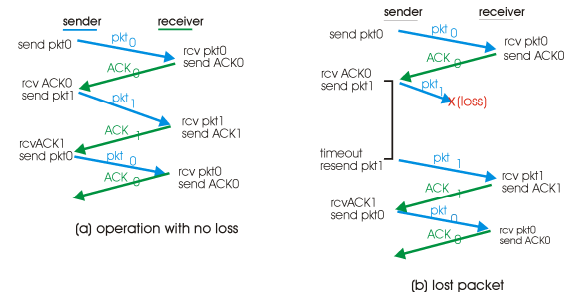
Transport Layer 10

rdt2.2: sender, receiver fragments



Transport Layer 8

rdt3.0 in action



Transport Layer 11

rdt3.0: channels with errors and loss

New assumption:

underlying channel can also lose packets (data or ACKs)

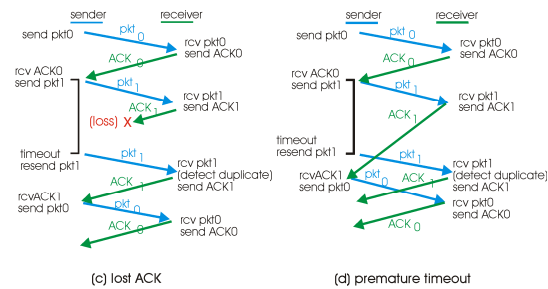
- checksum, seq. #, ACKs, retransmissions will be of help, but not enough

Approach: sender waits "reasonable" amount of time for ACK

- retransmits if no ACK received in this time
- if pkt (or ACK) just delayed (not lost):
 - retransmission will be duplicate, but use of seq. #'s already handles this
 - receiver must specify seq # of pkt being ACKed
- requires countdown timer

Transport Layer 9

rdt3.0 in action



Transport Layer 12

Performance of rdt3.0

- rdt3.0 works, but performance stinks
- example: 1 Gbps link, 15 ms e-e prop. delay, 1KB packet:

$$T_{\text{transmit}} = \frac{L \text{ (packet length in bits)}}{R \text{ (transmission rate, bps)}} = \frac{8\text{kb}/\text{pkt}}{10^{10} \text{ b/sec}} = 8 \text{ microsec}$$


- U_{sender} : utilization - fraction of time sender busy sending

$$U_{\text{sender}} = \frac{L/R}{RTT + L/R} = \frac{.008}{30.008} = 0.00027$$

- 1KB pkt every 30 msec → 33kB/sec thrupt over 1 Gbps link
- network protocol limits use of physical resources!

Transport Layer 13


Pipelining: increased utilization



$$U_{\text{sender}} = \frac{3 * L/R}{RTT + L/R} = \frac{.024}{30.008} = 0.0008$$

Transport Layer 16

rdt3.0: stop-and-wait operation




$$U_{\text{sender}} = \frac{L/R}{RTT + L/R} = \frac{.008}{30.008} = 0.00027$$

Transport Layer 14

Go-Back-N

Sender:

- k-bit seq # in pkt header
- "window" of up to N, consecutive unack'd pkts allowed




- ACK(n): ACKs all pkts up to, including seq # n - "cumulative ACK"
 - may receive duplicate ACKs (see receiver)
- timer for each in-flight pkt
- timeout(n): retransmit pkt n and all higher seq # pkts in window

Transport Layer 17

Pipelined protocols

Pipelining: sender allows multiple, "in-flight", yet-to-be-acknowledged pkts

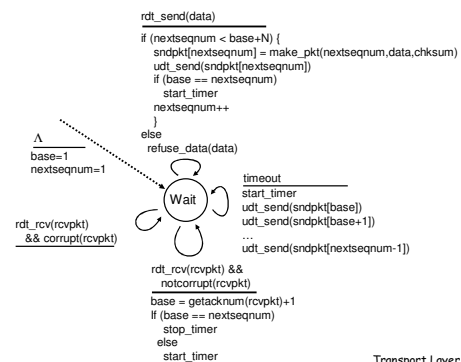
- range of sequence numbers must be increased
- buffering at sender and/or receiver



- Two generic forms of pipelined protocols: *go-Back-N*, *selective repeat*

Transport Layer 15

GBN: sender extended FSM



Transport Layer 18

GBN: receiver

- ACK-only: always send ACK for correctly-received pkt with highest *in-order* seq #
- may generate duplicate ACKs
 - need only remember `expectedseqnum`
- out-of-order pkt:
- discard (don't buffer) -> **no receiver buffering!**
 - Re-ACK pkt with highest in-order seq #

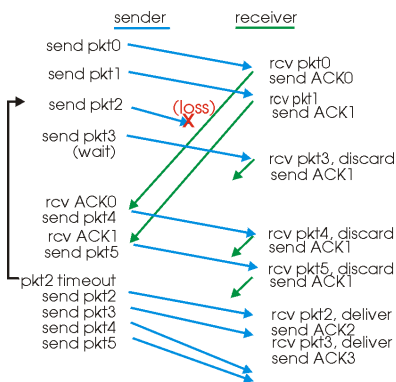
Transport Layer 19

Selective Repeat

- GBN is somewhat wasteful, as correctly received but out-of-order packets are sent again.
- In *selective repeat*, receiver *individually* acknowledges all correctly received pkts
 - buffers pkts, as needed, for eventual in-order delivery to upper layer
- sender only resends pkts for which ACK not received
 - sender timer for each unACKed pkt
- sender window
 - N consecutive seq #'s
 - again limits seq #'s of sent, unACKed pkts

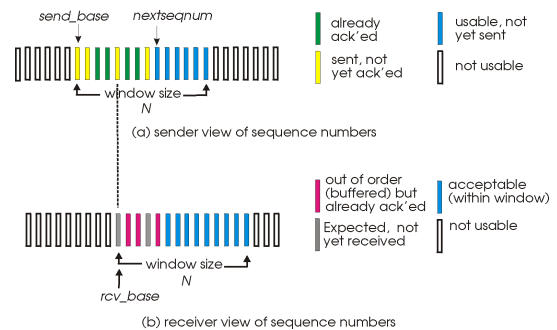
Transport Layer 22

GBN in action



Transport Layer 20

Selective repeat: sender, receiver windows



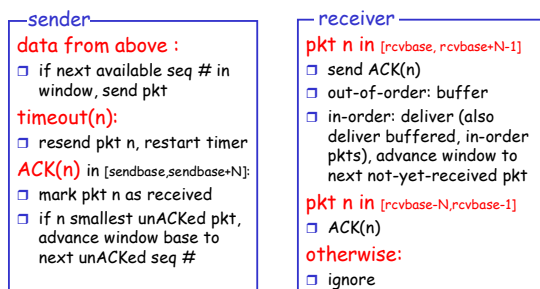
Transport Layer 23

Window Size

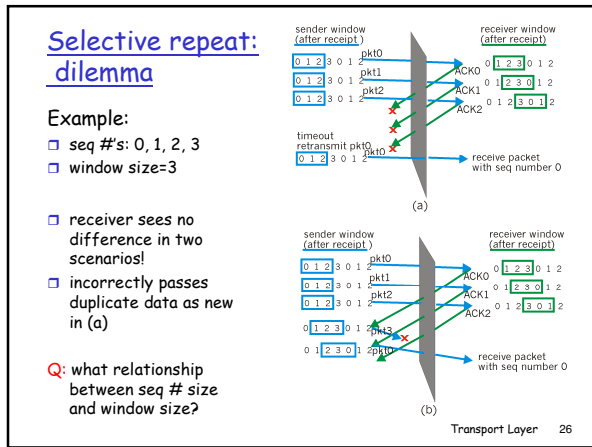
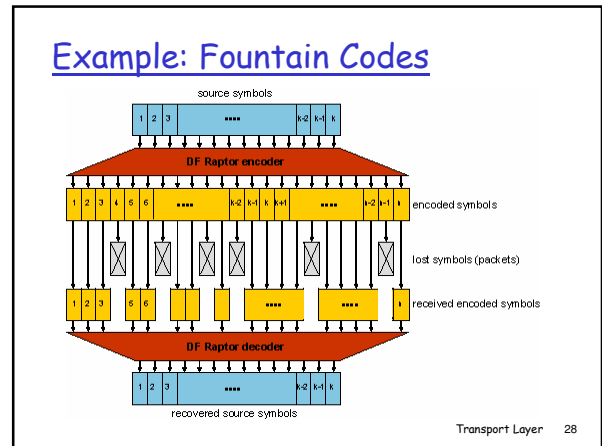
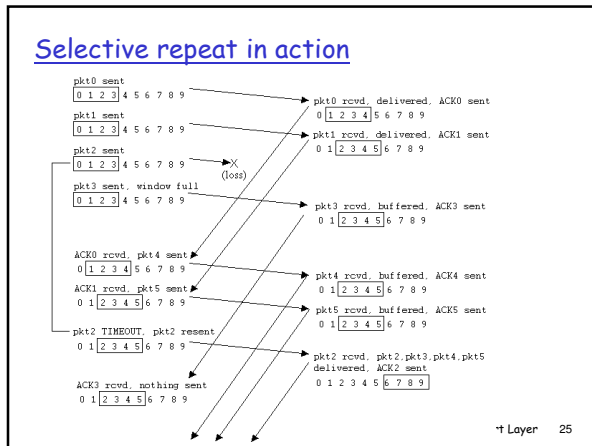
- The window size limits actual throughput.
- Why bother with having a window at all?
- It is used to do flow control and congestion control.
- This will be discussed next week.

Transport Layer 21

Selective repeat



Transport Layer 24



Forward Erasure/Error Correction: A Different Approach to RDT

- Our approach to reliable data delivery is based on ACKs and retransmissions, i.e. **feedback**.
- Long RTTs => long delays and/or low throughput
- An alternative approach is via **forward** corrections for errors and losses.

Transport Layer 27