

Last time

- Reliable Data Transfer
 - ◆ Provide rdt over unreliable network layer
 - ◆ FSM model
 - ◆ rdt 1.0: rdt over reliable channels
 - ◆ rdt 2.0: rdt over channels with bit errors
 - ◆ rdt 2.1: handle garbled ACKs/NAKs
 - ◆ rdt 2.2: remove need for NAKs

- Midterm review

This time

- Reliable Data Transfer with packet loss
- Pipelining

rdt3.0: channels with errors *and* loss

New assumption:

underlying channel can also lose packets (data or ACKs)

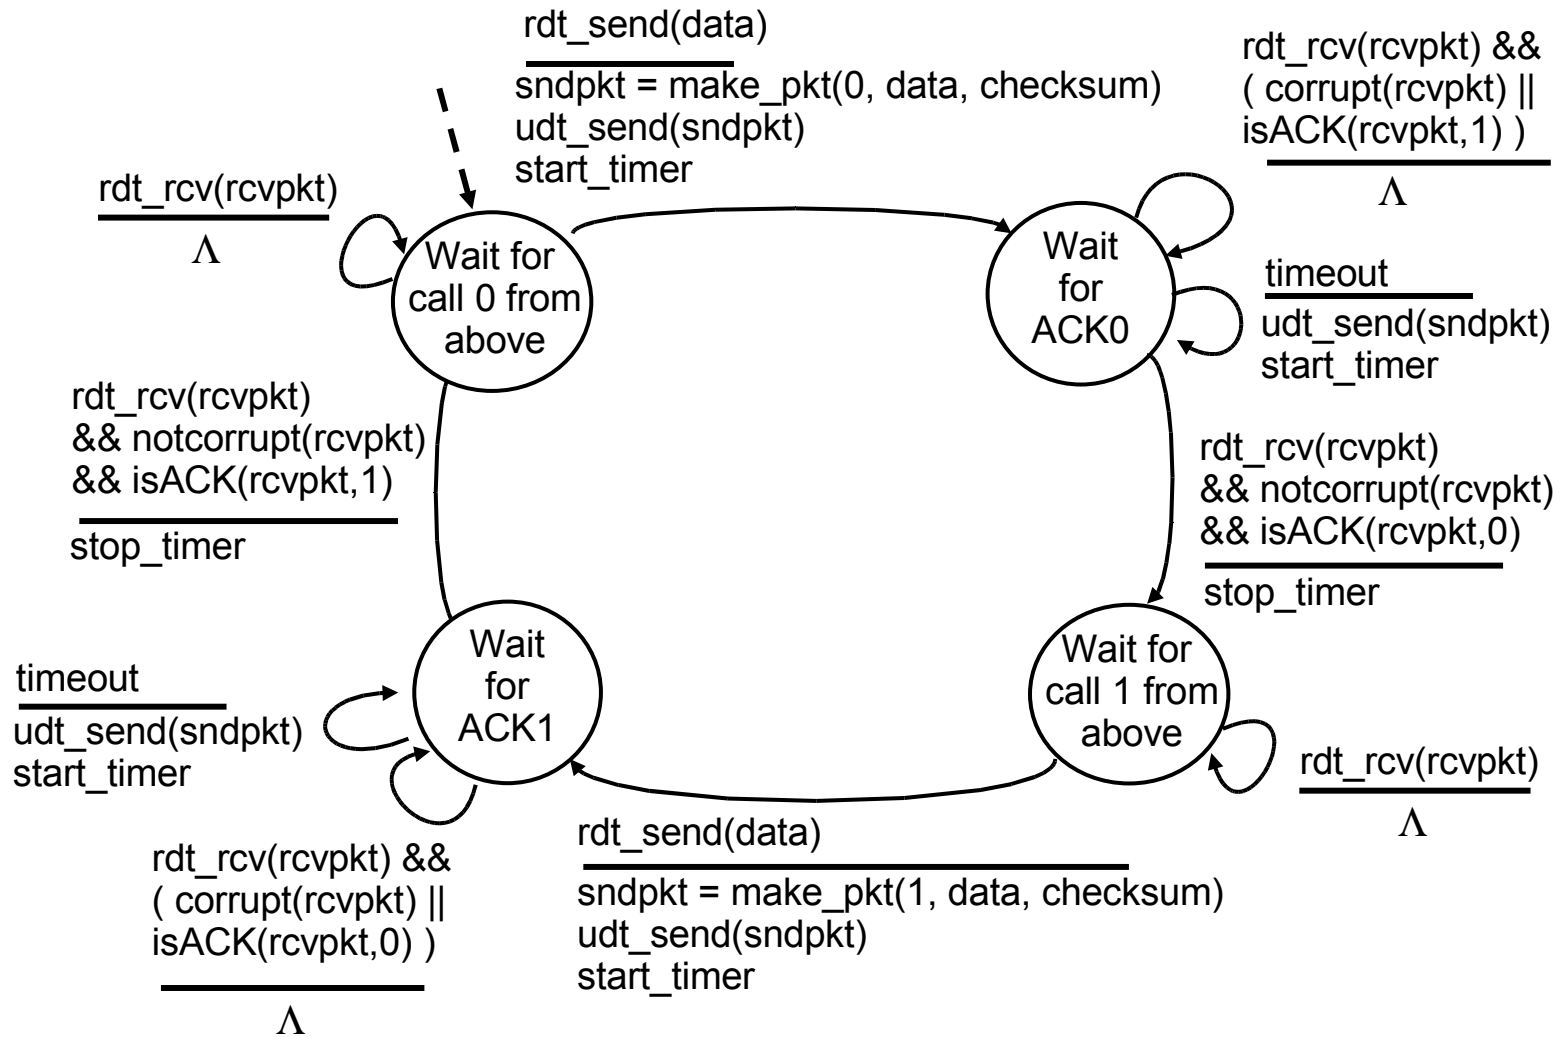
- ◆ checksum, seq. #, ACKs, retransmissions will be of help, but not enough

Approach: sender waits

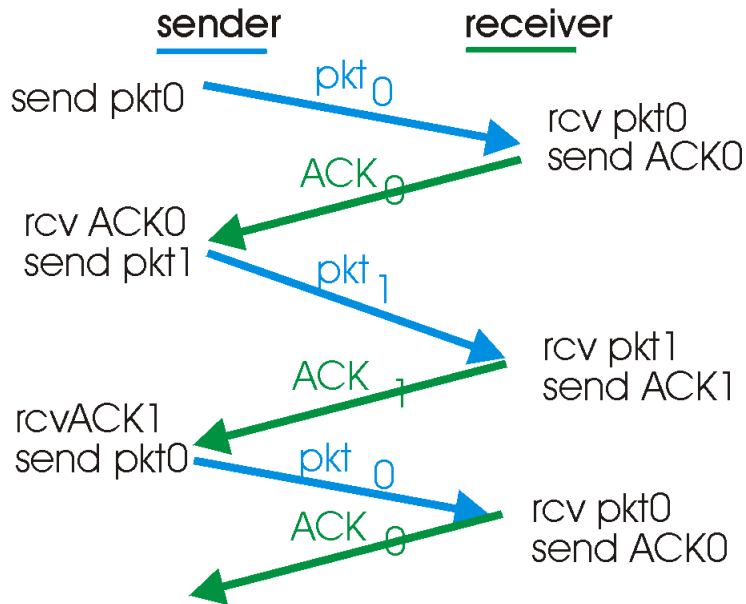
“reasonable” amount of time for ACK

- 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

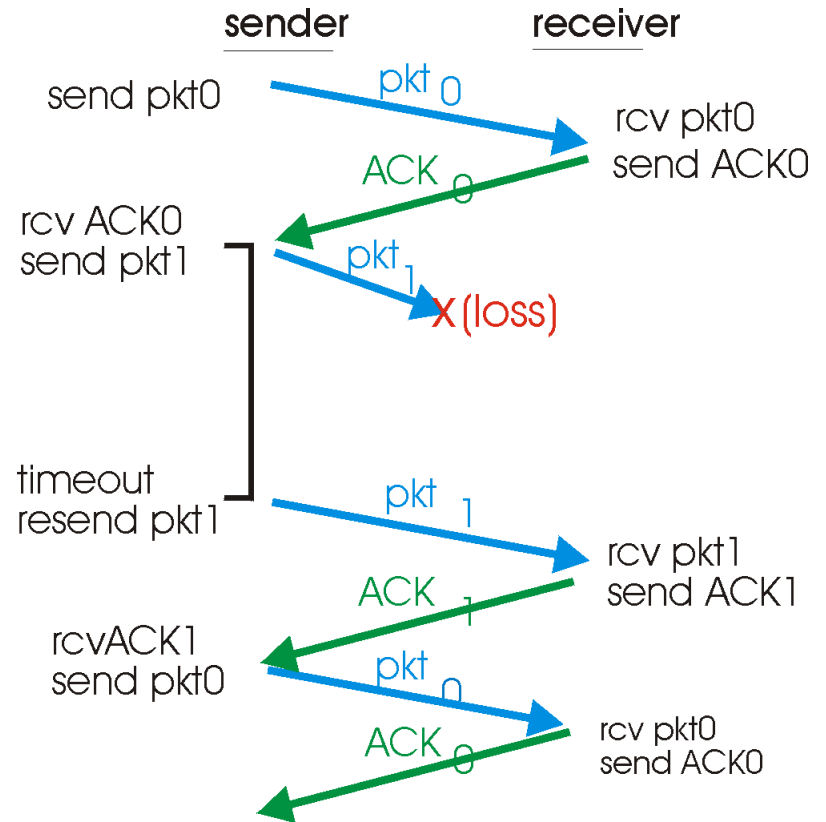
rdt3.0 sender



rdt3.0 in action

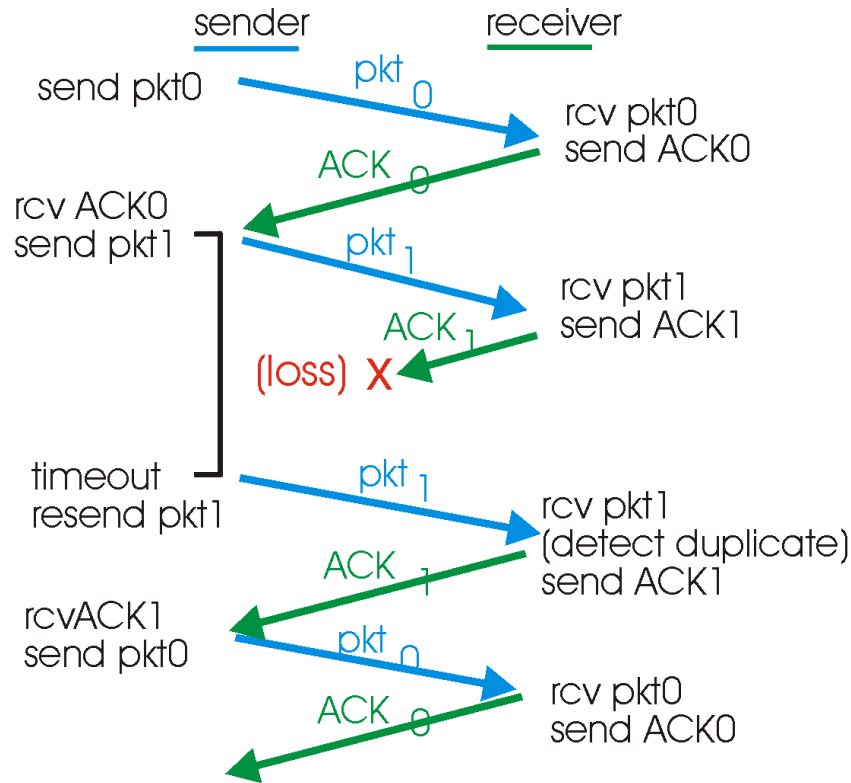


(a) operation with no loss

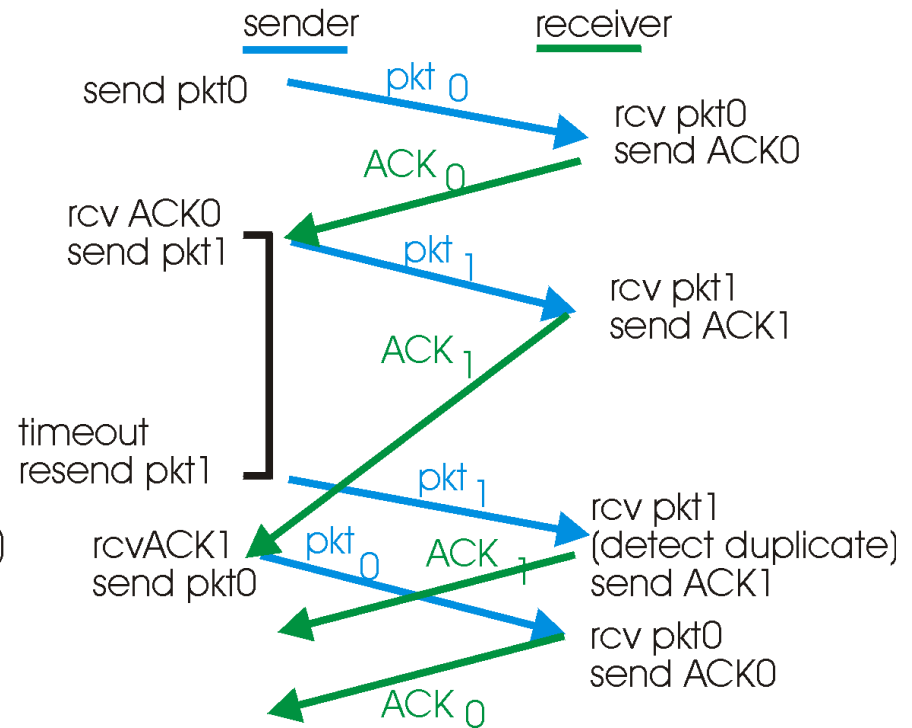


(b) lost packet

rdt3.0 in action



(c) lost ACK



(d) premature timeout

Performance of rdt3.0

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


$$T_{\text{transmit}} = \frac{L \text{ (packet length in bits)}}{R \text{ (transmission rate, bps)}} = \frac{8\text{kb/pkt}}{10^{**9} \text{ b/sec}} = 8 \text{ microsec}$$

- ◆ 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 throughput over 1 Gbps link
- ◆ network protocol limits use of physical resources!

rdt3.0: stop-and-wait operation




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

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




(a) a stop-and-wait protocol in operation

(b) a pipelined protocol in operation

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

Pipelining: increased utilization




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

Increase utilization by a factor of 3!

Go-Back-N

Sender:

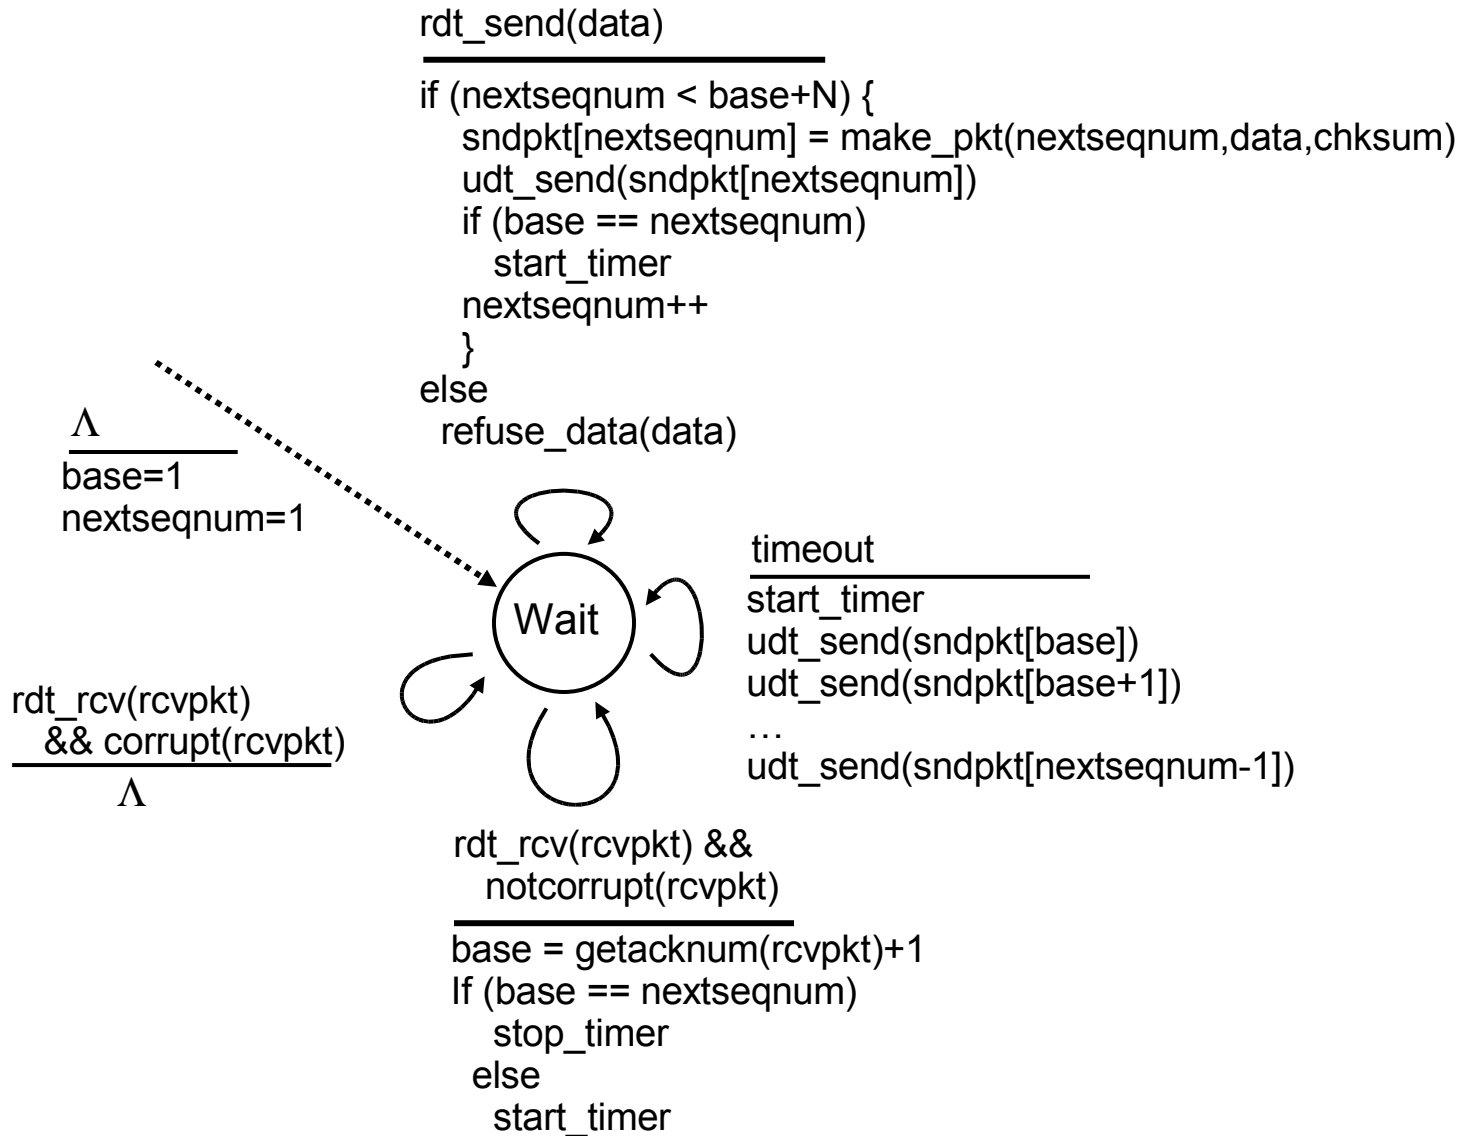
- k-bit seq # in pkt header
- “window” of up to N, consecutive unack’ed pkts allowed



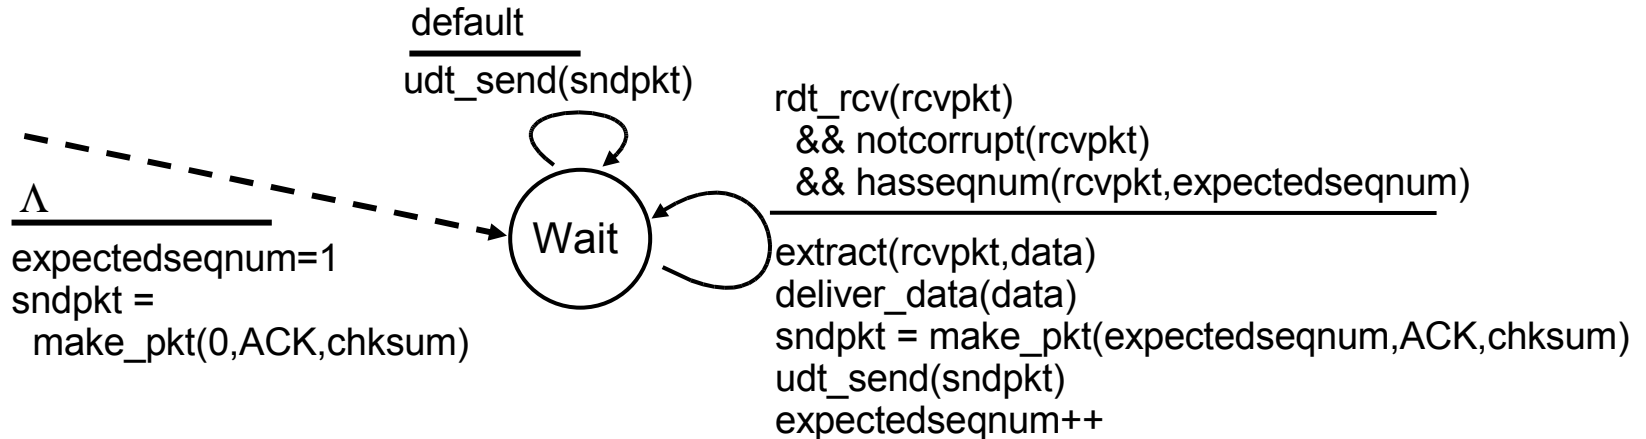
- ACK(n): ACKs all pkts up to, including seq # n - “cumulative ACK”
 - ◆ may receive duplicate ACKs (see receiver)
- timer for each in-flight pkt
- *timeout(n)*: retransmit pkt n and all higher seq # pkts in window

See the applet on UW-ACE!

GBN: sender extended FSM



GBN: receiver extended FSM



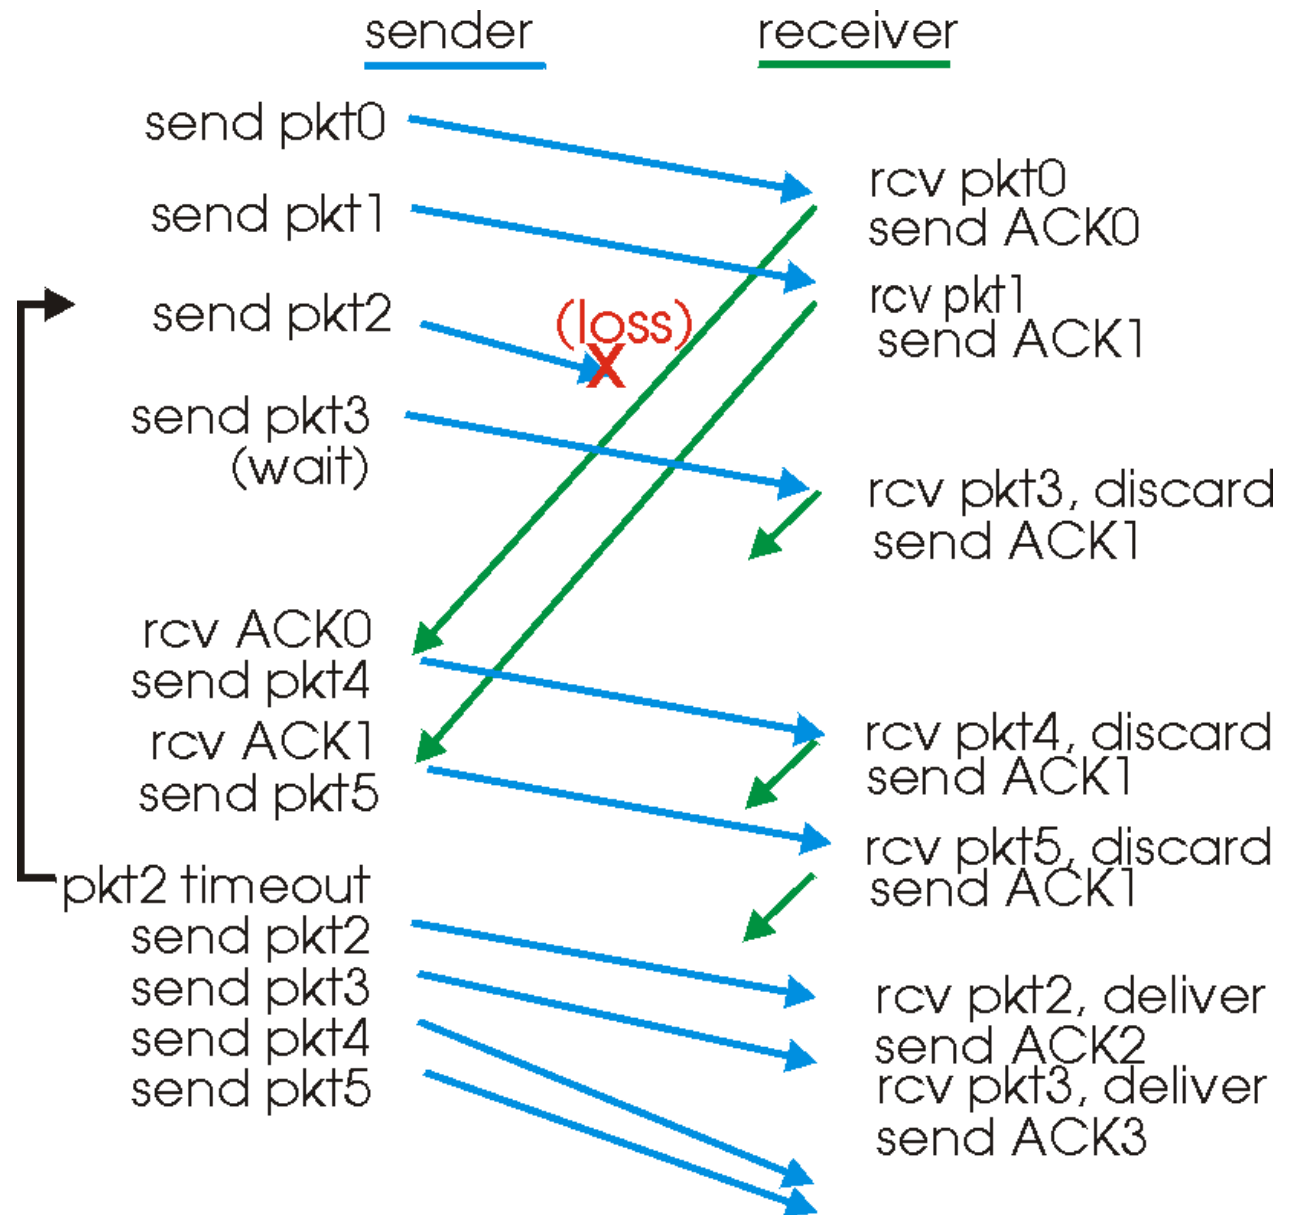
ACK-only: always send ACK for correctly-received pkt with highest *in-order* seq #

- ◆ 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 #

GBN in action




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

See the applet on UW-ACE!

Selective repeat: sender, receiver windows



(a) sender view of sequence numbers



(b) receiver view of sequence numbers

Selective repeat

sender

data from above :

- if next available seq # in window, send pkt

timeout(n):

- resend pkt n, restart timer

ACK(n) in [sendbase, sendbase+N]:

- mark pkt n as received
- if n smallest unACKed pkt, advance window base to next unACKed seq #

receiver

pkt n in [rcvbase, rcvbase+N-1]

- send ACK(n)
- out-of-order: buffer
- in-order: deliver (also deliver buffered, in-order pkts), advance window to next not-yet-received pkt

pkt n in [rcvbase-N, rcvbase-1]

- ACK(n)

otherwise:

- ignore



Selective repeat: dilemma

Example:

- seq #'s: 0, 1, 2, 3
- window size=3

- receiver sees no difference in two scenarios!
- incorrectly passes duplicate data as new in (a)

- Q: what relationship between seq # size and window size?



Recap

- Reliable Data Transfer with packet loss
 - ◆ rdt3.0

- Pipelining
 - ◆ Problems with stop-and-wait
 - ◆ Go-Back-N
 - ◆ Selective-Repeat

Next time

- UDP socket programming
- TCP