

CSCI 491-01

Topics: Internet Programming

Fall 2008

Transport Layer

Derek Leonard
Hendrix College

October 6, 2008

Original slides copyright © 1996-2007 J.F Kurose and K.W. Ross

Chapter 3: Roadmap

3.1 Transport-layer services

3.2 Multiplexing and demultiplexing

3.3 Connectionless transport: UDP

3.4 Principles of reliable data transfer

3.5 Connection-oriented transport: TCP

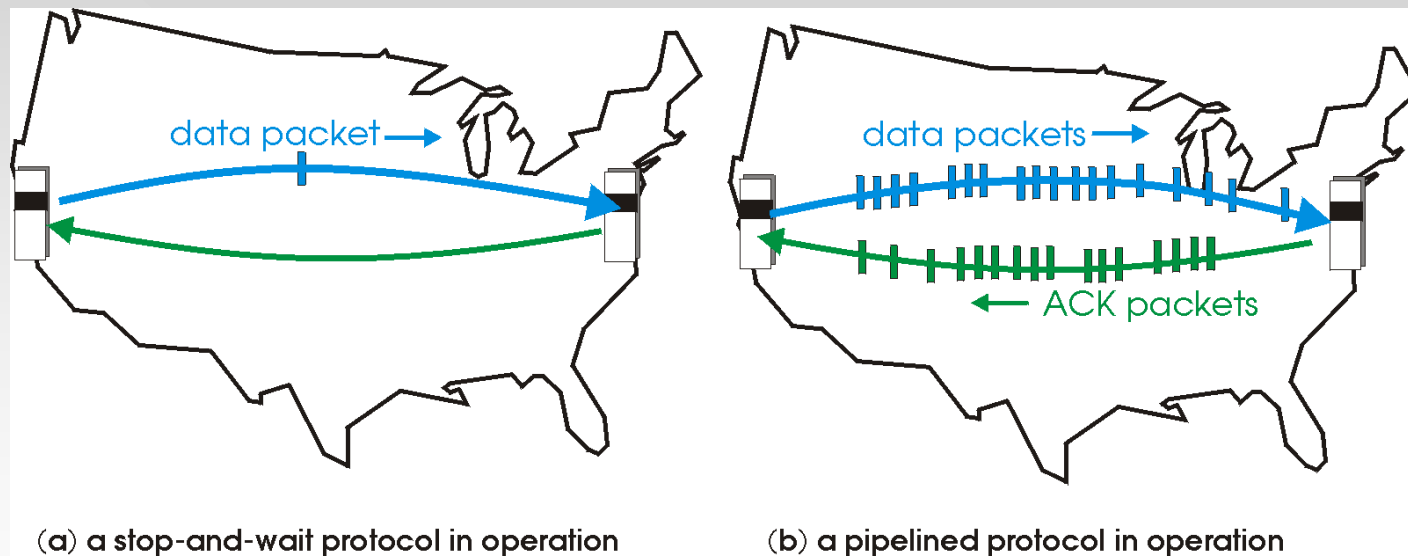
- Segment structure
- Reliable data transfer
- Flow control
- Connection management

3.6 Principles of congestion control

3.7 TCP congestion control


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* and *Selective Repeat*

Pipelining: Increased Utilization




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

Go-back-N

Sender:

- **Window** of up to N consecutive unack'ed pkts allowed
- k -bit seq numbers in pkt header



- ACK(n): ACKs all pkts up to, including seq # n – called “**cumulative ACK**”
 - May generate duplicate ACKs (see receiver)
- Timer for the oldest unacknowledged pkt n :
 - *Timeout(n)*: retransmit pkt n and all higher seq # pkts in current window

GBN: Sender Extended FSM

```

rdt_send(data)
  if (nextseqnum < base+N) {
    sndpkt[nextseqnum] = make_pkt(nextseqnum,data,chksum)
    udt_send (sndpkt[nextseqnum])
    if (base == nextseqnum) start_timer
    nextseqnum++
  }
else refuse_data(data)

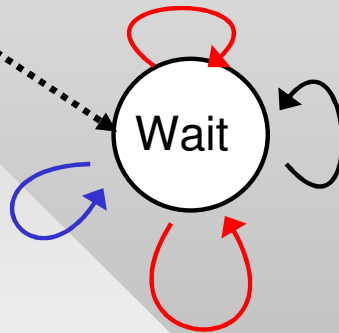
```

Λ
 base=1
 nextseqnum=1

rdt_rcv(rcvpkt)
 && corrupt(rcvpkt)

Λ

else refuse_data(data)



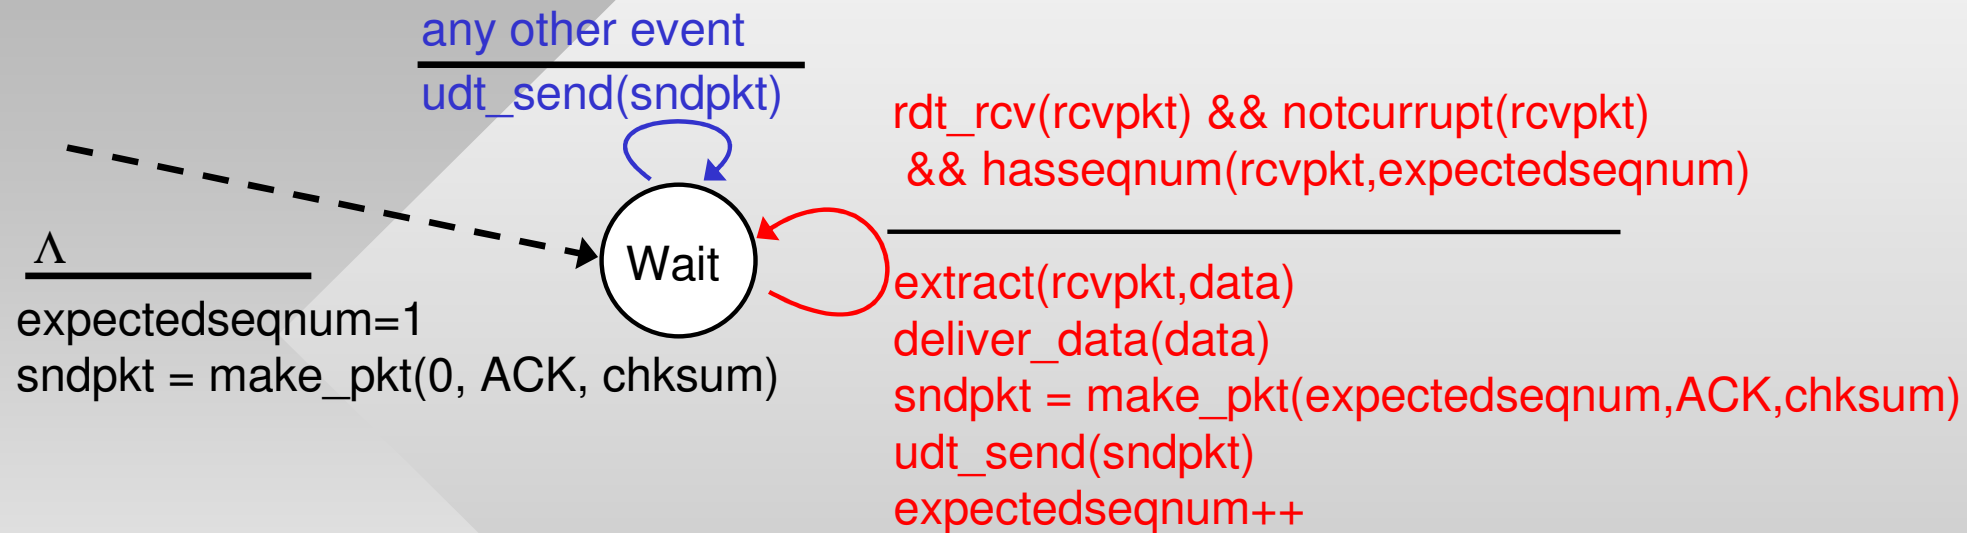
timeout
 start_timer
 udt_send(sndpkt[base])
 udt_send(sndpkt[base+1])
 ...
 udt_send(sndpkt[nextseqnum-1])

```

rdt_rcv(rcvpkt) &&
notcorrupt(rcvpkt)
  tmp = getacknum(rcvpkt)+1
  if (tmp > base) {    base = tmp;
    if (base == nextseqnum)
      stop_timer // last ACK in window
    else start_timer
  }

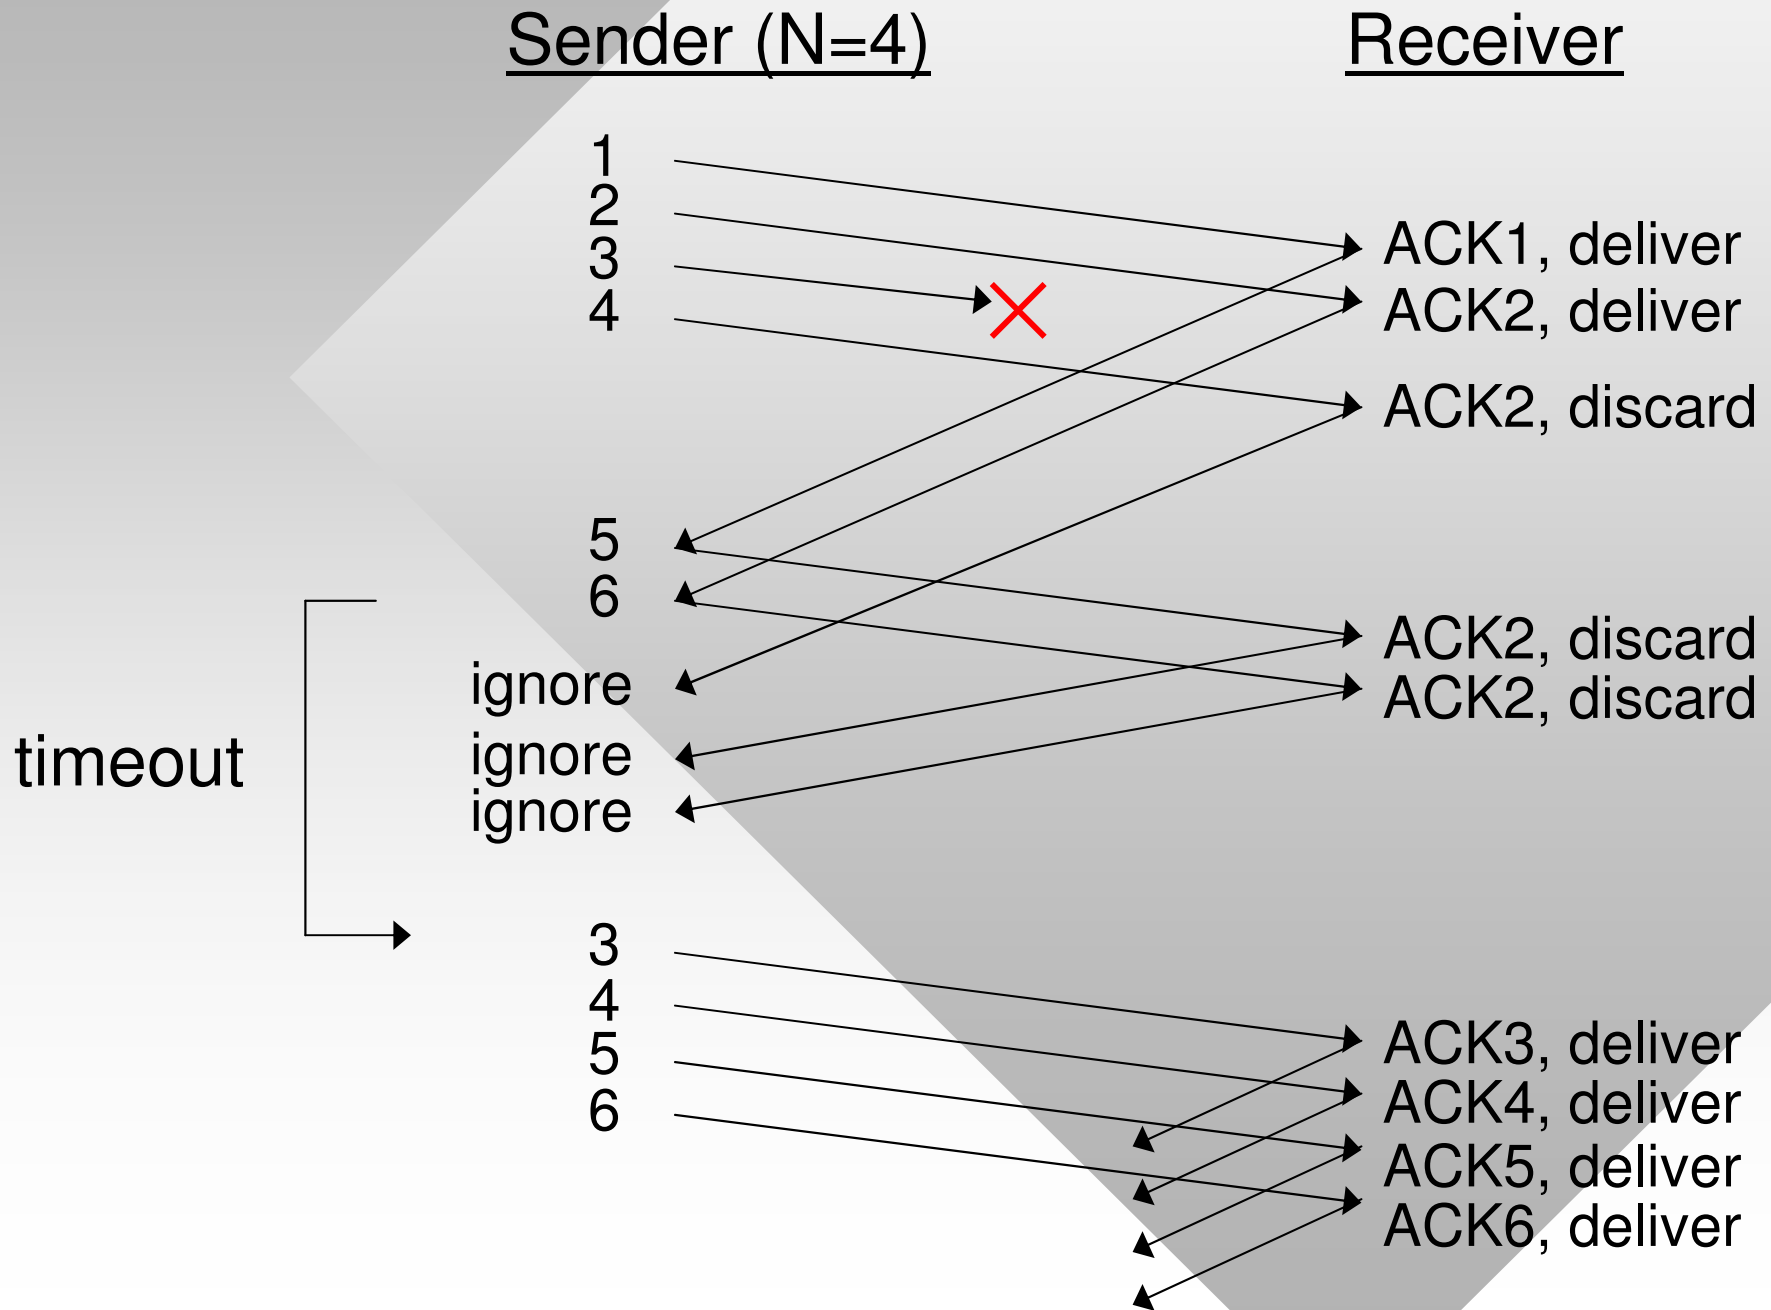
```

GBN: Receiver Extended FSM



- ACK-only: always send ACK for correctly-received pkt with highest *in-order* seq #
 - Duplicate ACKs during loss
 - Need only remember **expectedseqnum**
- Out-of-order pkt:
 - Discard → **no receiver buffering!**
 - Re-ACK pkt with highest in-order seq #


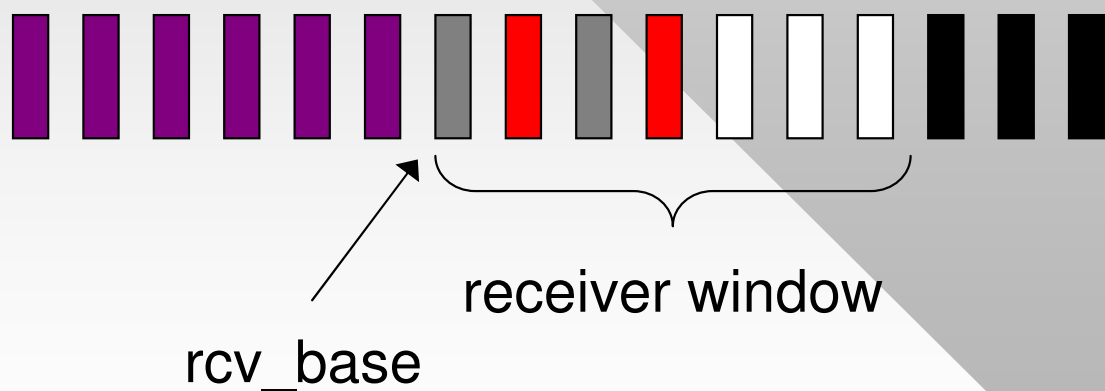

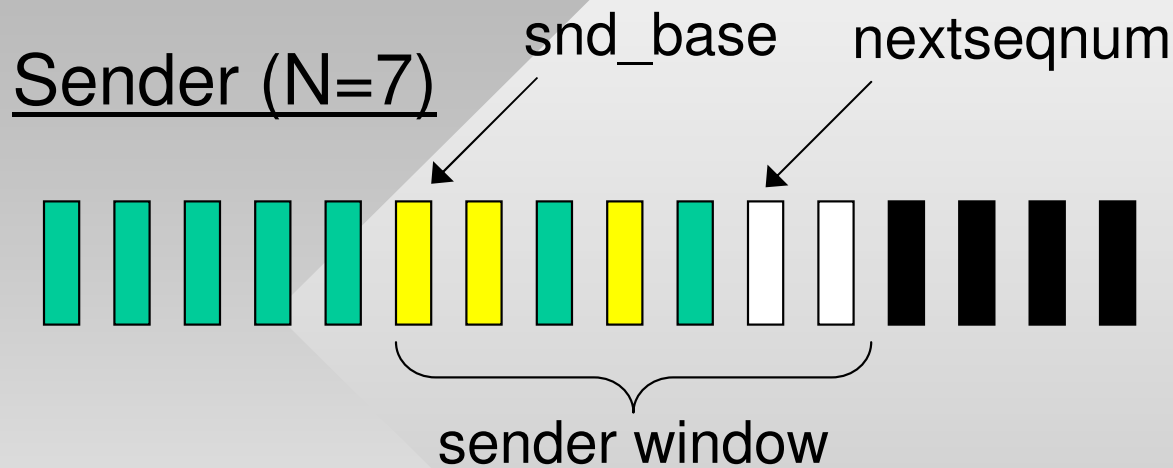
GBN in Action



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
 - **Separate timer for each unACKed pkt**
- Sender window
 - N consecutive packets in $[\text{snd_base}, \text{snd_base}+N-1]$

Selective Repeat: Sender, Receiver Windows



Selective Repeat

sender

Data from above :

- If next available seq # in window, send pkt

Timeout(n):

- Resend pkt n, restart timer n

ACK(n) in [sendbase, sendbase+N-1]:

- Mark pkt n as received
- If n smallest unACKed pkt, advance window base to next unACKed seq #

receiver

- Receive pkt n in [rcvbase, rcvbase+N-1]
- Send ACK(n)
- Out-of-order: buffer
- In-order: deliver, advance rcvbase to next not-yet-received pkt, deliver all buffered, in-order pkts

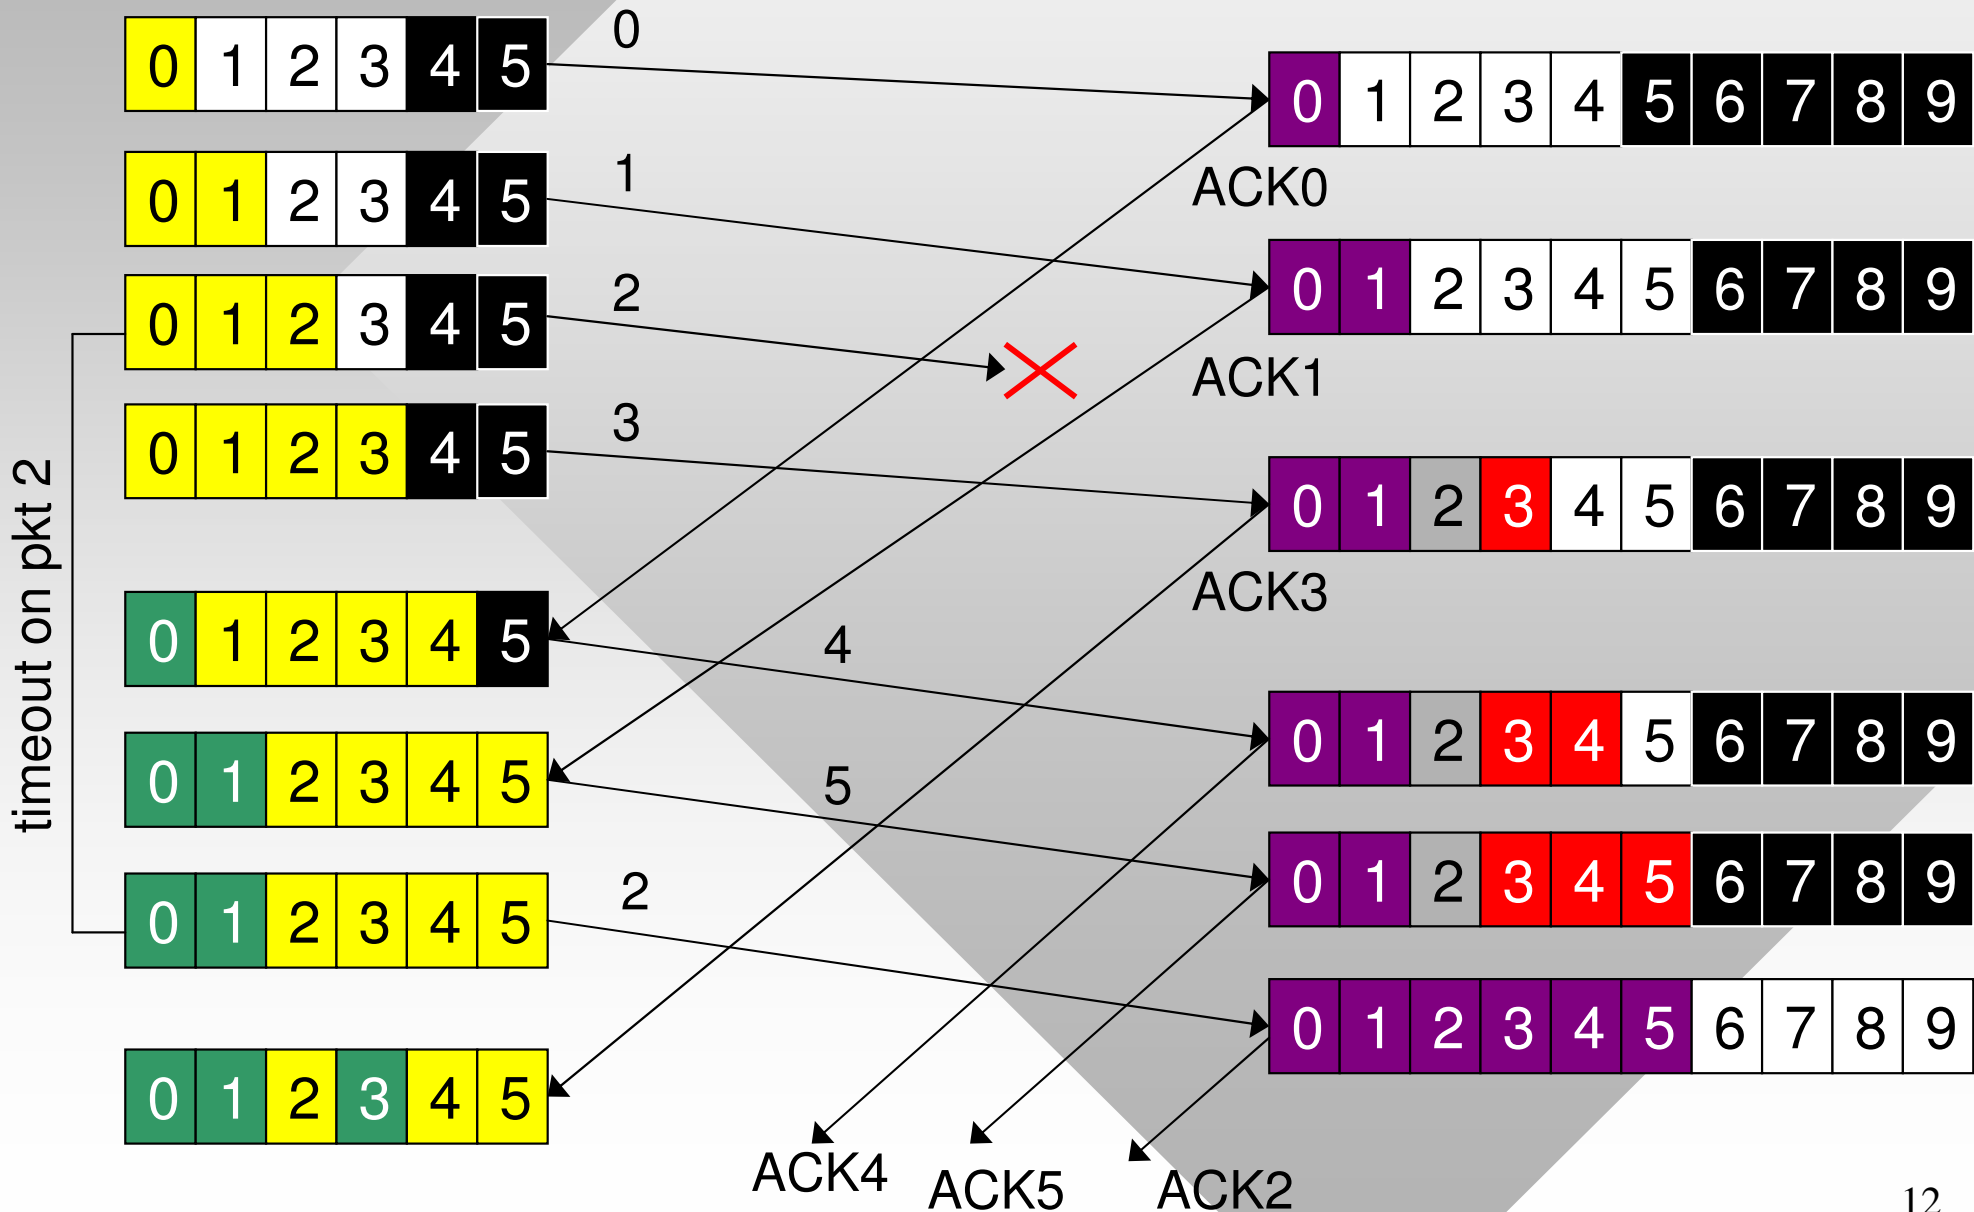
Pkt n in [rcvbase-N,rcvbase-1]

- ACK(n)

Otherwise:

- Ignore

Selective Repeat in Action (N=4)



Selective Repeat: Dilemma

Q: How many distinct seq #s are needed for window size N in selective repeat?

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)

