

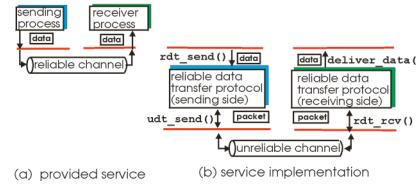
Reliable Data Transfer

Transport Layer 3-1

Principles of Reliable data transfer

- important in app., transport, link layers

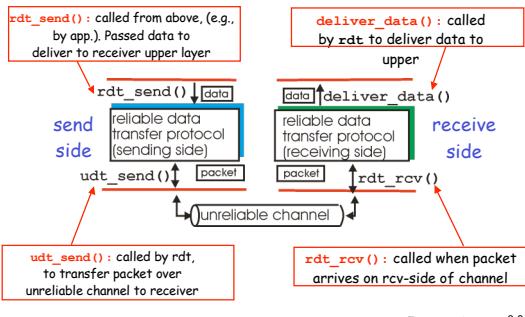
- top-10 list of important networking topics!



- characteristics of unreliable channel will determine complexity of reliable data transfer protocol (rdt)

Transport Layer 3-2

Reliable data transfer: getting started



Transport Layer 3-3

Reliable data transfer: getting started

We'll:

- incrementally develop sender, receiver sides of reliable data transfer protocol (rdt)

- consider only unidirectional data transfer
 - but control info will flow on both directions!

- use finite state machines (FSM) to specify sender, receiver

state: when in this "state" next state uniquely determined by next event
event causing state transition
actions taken on state transition

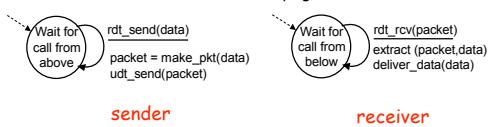


Transport Layer 3-4

Rdt1.0: reliable transfer over a reliable channel

- underlying channel perfectly reliable
 - no bit errors
 - no loss of packets

- separate FSMs for sender, receiver:
 - sender sends data into underlying channel
 - receiver reads data from underlying channel



Transport Layer 3-5

Rdt2.0: channel with bit errors

- underlying channel may flip bits in packet
 - checksum to detect bit errors

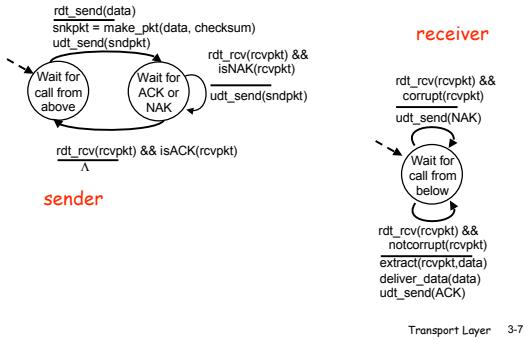
- the question: how to recover from errors:

- acknowledgements (ACKs):** receiver explicitly tells sender that pkt received OK
- negative acknowledgements (NAKs):** receiver explicitly tells sender that pkt had errors
- sender retransmits pkt on receipt of NAK

- new mechanisms in rdt2.0 (beyond rdt1.0):
 - error detection
 - receiver feedback: control msgs (ACK, NAK) rcvr->sender

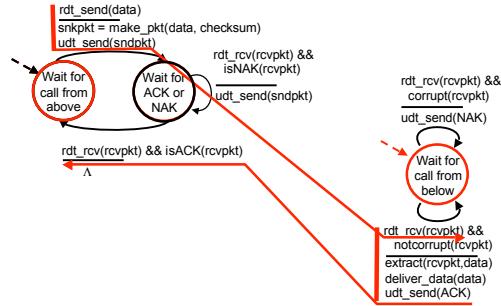
Transport Layer 3-6

rdt2.0: FSM specification



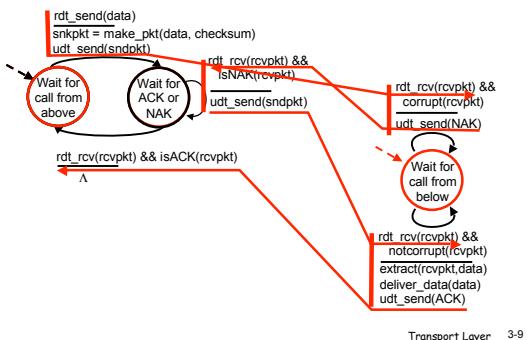
Transport Layer 3-7

rdt2.0: operation with no errors



Transport Layer 3-8

rdt2.0: error scenario



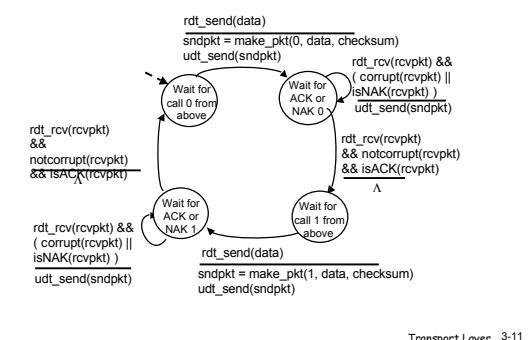
Transport Layer 3-9

rdt2.0 has a fatal flaw!

- **Stop and Wait**
 - Sender sends one packet, then waits for receiver response
- **What happens if ACK/NAK corrupted?**
 - sender doesn't know what happened at receiver!
 - can't just retransmit: possible duplicate
- **Handling duplicates:**
 - sender adds *sequence number* to each pkt
 - sender retransmits current pkt if ACK/NAK garbled
 - receiver discards (doesn't deliver up) duplicate pkt

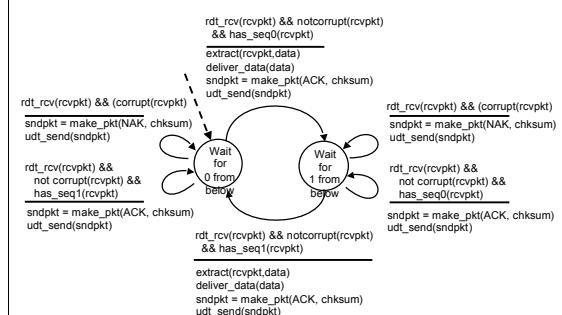
Transport Layer 3-10

rdt2.1: sender, handles garbled ACK/NAKs



Transport Layer 3-11

rdt2.1: receiver, handles garbled ACK/NAKs



Transport Layer 3-12

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

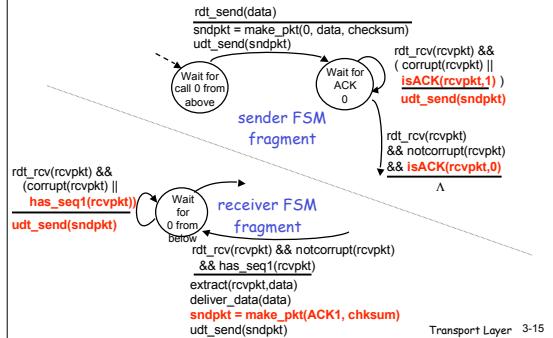
Transport Layer 3-13

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 3-14

rdt2.2: sender, receiver fragments



Transport Layer 3-15

rdt3.0: channels with errors and loss

New assumption:

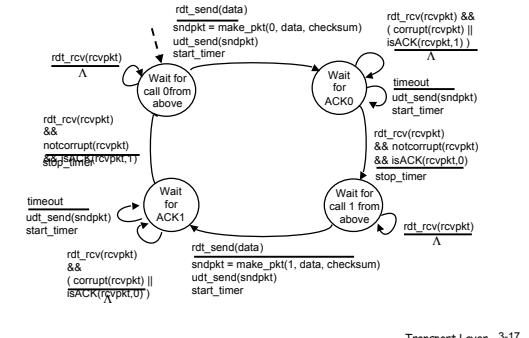
- underlying channel can also lose packets (data or ACKs)
- checksum, seq. #, ACKs, retransmissions will be of help, but not enough

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

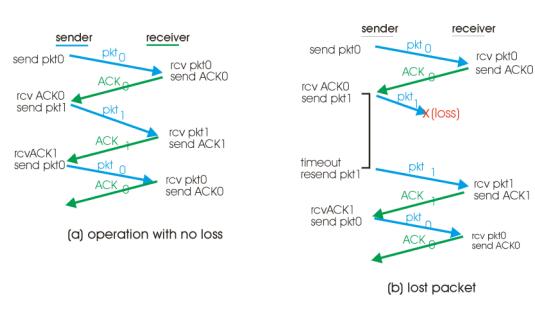
Transport Layer 3-16

rdt3.0 sender

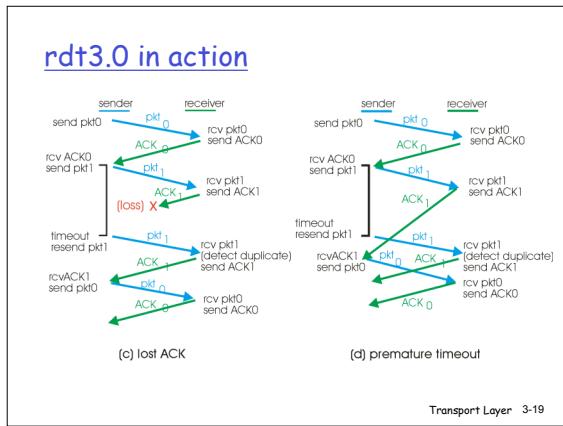


Transport Layer 3-17

rdt3.0 in action



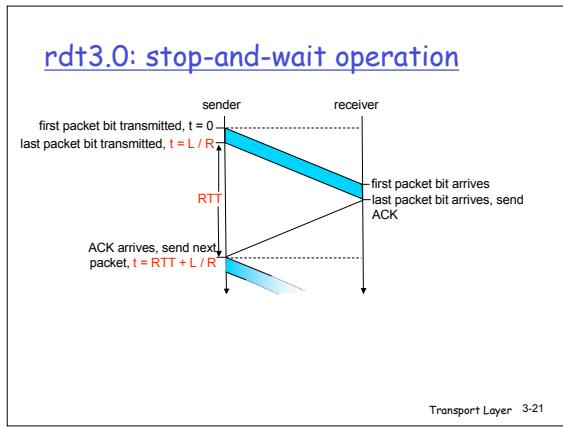
Transport Layer 3-18



Performance of rdt3.0

- ❑ rdt3.0 works, but performance stinks
 - ❑ Why?

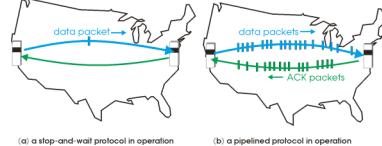
Transport Layer 3-20



Pipelined protocols

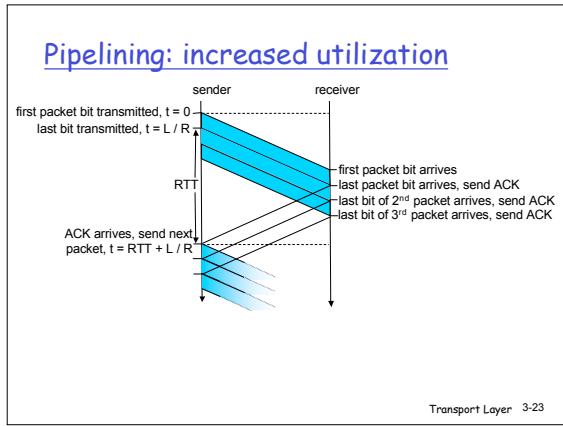
Pipelining: sender allows multiple, "in-flight", yet-to-be-acknowledged pkts

- 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 3-22



Go-Back-N

Sender:

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

already ack'd	sent, not yet ack'd	not usable
---------------	---------------------	------------
 - ❑ ACK(n): ACKs all pkts up to, including seq # n - "cumulative ACK"
 - may deceive duplicate ACKs (see receiver)
 - ❑ timer for each in-flight pkt
 - ❑ timeout(n): retransmit pkt n and all higher seq # pkts in window

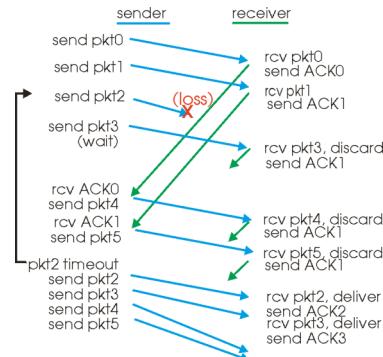
Transport Layer 3-24

GBN: receiver

- ACK-only:** always send ACK for correctly-received pkt with highest *in-order* seq #
- o may generate duplicate ACKs
- o need only remember **expectedseqnum**
- out-of-order pkt:**
 - o discard (don't buffer) -> **isn't this bad?**
 - o Re-ACK pkt with highest in-order seq #

Transport Layer 3-25

GBN in action



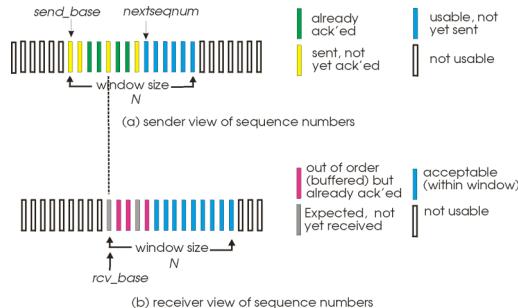
Transport Layer 3-26

Selective Repeat

- receiver individually acknowledges all correctly received pkts**
 - o buffers pkts, as needed, for eventual in-order delivery to upper layer
- sender only resends pkts for which ACK not received**
 - o sender timer for each unACKed pkt
- sender window**
 - o N consecutive seq #'s
 - o again limits seq #'s of sent, unACKed pkts

Transport Layer 3-27

Selective repeat: sender, receiver windows



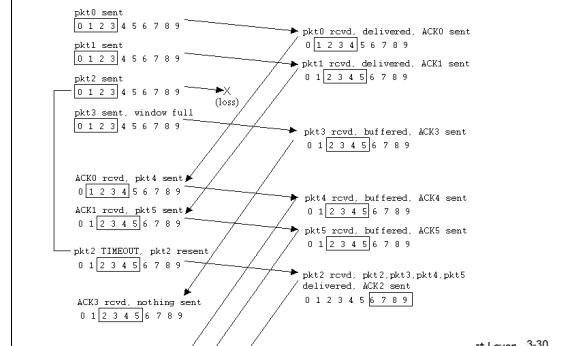
Transport Layer 3-28

Selective repeat

- | | |
|--|---|
| sender | receiver |
| data from above : | |
| <ul style="list-style-type: none"> □ if next available seq # in window, send pkt | <ul style="list-style-type: none"> pkt n in [rcvbase, rcvbase+N-1] □ send ACK(n) |
| timeout(n): | <ul style="list-style-type: none"> □ out-of-order: buffer |
| <ul style="list-style-type: none"> □ resend pkt n, restart timer | <ul style="list-style-type: none"> □ in-order: deliver (also deliver buffered, in-order pkts), advance window to next not-yet-received pkt |
| ACK(n) in [sendbase, sendbase+N]: | <ul style="list-style-type: none"> pkt n in [rcvbase-N, rcvbase-1] □ ACK(n) |
| <ul style="list-style-type: none"> □ if n smallest unACKed pkt, advance window base to next unACKed seq # | <ul style="list-style-type: none"> otherwise: □ ignore |

Transport Layer 3-29

Selective repeat in action

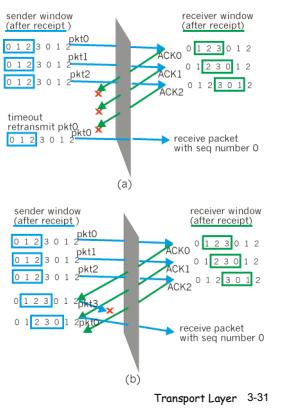


† Layer 3-30

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?



Transport Layer 3-31