Lecture 10: Transport Layer Reliable Data Transfer and Seq #s

COMP 332, Spring 2024 Victoria Manfredi

WESLEYAN UNIVERSITY



Acknowledgements: materials adapted from Computer Networking: A Top Down Approach 7th edition: ©1996-2016, J.F Kurose and K.W. Ross, All Rights Reserved as well as from slides by Abraham Matta at Boston University, and some material from Computer Networks by Tannenbaum and Wetherall.

Today

Announcements

- homework 4 due Thursday, 11:59p

Recap

- reliable data transport over channels with errors and loss

Pipelined protocols

- go-back-N
- selective repeat
- sequence numbers in practice

TCP overview

Reliable Data Transport REVIEW

Summary of techniques and uses



Will see: # of seq #s must be > 2x window size if reordering

Are we done?

Have we solved reliable communication over an unreliable channel?

Reliable Data Transport PIPELINED PROTOCOLS

rdt3.0: stop-and-wait operation



Problem: how to maintain high link utilization?

Get rid of stop-and wait

Use pipelining (aka sliding-window protocols), like in HTTP

- sender allows multiple, in-flight, yet-to-be-acknowledged pkts
 - send up to N packets at a time, unacked
 - range of seq #s must be increased
 - sender needs more memory to buffer outstanding unacked packets



Achieves higher link utilization than stop-and-wait!

Increased utilization with pipelining



Pipelined protocols

Send N packets without receiving ACKs. How to ACK now?

Cumulative ACKs: Go-Back-N protocol

– sender

- has timer for oldest unacked pkt
- when timer expires: retransmit all unacked pkts
- pkts received correctly may be retransmitted
- receiver only sends cumulative ack, doesn't ack pkt if gap

Selective ACKs: Selective Repeat protocol

– sender

- has timer for each unacked pkt
- when timer expires, retransmit only unacked pkt
- only corrupted/lost pkts are retransmitted
- receiver sends individual ack for each pkt

How pipelining protocols work

What is window size on stop and wait protocol?

Use sliding window

- how sender keeps track of what it can send
- window: set of N adjacent seq #s
 - only send packets in window



If window large enough, will fully utilize link

Pipelined Protocols GO-BACK-N

Go-Back-N: sender

Window of up to N consecutive unacked pkts allowed

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



Go-Back-N: sender FSM



Go-Back-N: receiver FSM

Out-of-order pkt and all other cases



Correct pkt with highest in-order seq #

- send ACK, may be duplicate ACK
- need only remember expectedseqnum



rdt_rcv(rcvpkt) && notcorrupt(rcvpkt) && hasseqnum(rcvpkt,expectedseqnum)

extract(rcvpkt,data) deliver_data(data) sndpkt = make_pkt(expectedseqnum,ACK,chksum) udt_send(sndpkt) expectedseqnum++

Retransmit windowsize worth of packets for 1 error large window size ⇒ large delays

Go-Back-N in action

<u>sender window (l</u>	<u>N=4)</u> <u>sender</u>	<u>receiver</u>
<mark>0 1 2 3</mark> 4 5 6 7 8	send pkt0	
<mark>0 1 2 3</mark> 4 5 6 7 8	send pkt1	raceive pltto cond acko
<mark>0 1 2 3</mark> 4 5 6 7 8	send pkt2-	
<mark>0 1 2 3</mark> 4 5 6 7 8	send pkt3	Y loss receive pkt1, send ack1
	(wait)	receive pkt3, discard,
0 <mark>1 2 3 4 </mark> 5 6 7 8	rcv ack0, send pkt4	(Te)senu acki
0 1 <mark>2 3 4 5 </mark> 6 7 8	rcv ack1, send pkt5	receive pkt4, discard,
	ignore duplicate ACK	receive pkt5, discard,
	💮 pkt 2 timeout _	(re)send ack1
0 1 <mark>2 3 4 5 </mark> 6 7 8	send pkt2	
0 1 <mark>2 3 4 5</mark> 6 7 8	send pkt3	
0 1 <mark>2 3 4 5 </mark> 6 7 8	send pkt4	rcv pkt2, deliver, send ack2
0 1 <mark>2 3 4 5 </mark> 6 7 8	sena pkt5	rcv pkt3, deliver, send ack3
		rcy nkt5 deliver send ack5

Go-Back-N summary

Pros

- no receiver buffering
 - saves resources by requiring packets to arrive in-order
 - avoids large bursts of packet delivery to higher layers
- simpler buffering & protocol processing
 - can easily detect duplicates if out-of-sequence packet is received

Cons

- wastes capacity
 - on timeout for packet N sender retransmits from N all over again (all outstanding packets) including potentially correctly received packets

Tradeoff: buffering/processing complexity vs. capacity (time vs. space)

Pipelined Protocols SELECTIVE REPEAT

Selective repeat

Rather than ACK cumulatively, ACKs selectively

Receiver

- individually ACKs 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
- limits seq #s of sent, unACKed pkts

Selective repeat: sender, receiver windows



Selective repeat sender

Event: data from above

action: if has next available seq # in window, send packet, start timer

Event: timeout(n)

- action: resend packet n, restart timer

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

- action
 - mark packet n as received
 - if n is smallest unACKed packet
 - advance window base to next unACKed seq #

Selective repeat receiver

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

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

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

– action: send ACK(n)

Event: otherwise

- action: ignore

Selective repeat in action



Selective repeat: dilemma

Example

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



No problem...

Selective repeat: dilemma

Example

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



Problem: duplicate data accepted as new: receiver sees no difference in two scenarios!

Q: what is relationship between seq # size and window size to avoid problem in (b)?

Selective repeat summary

Q: When is selective repeat useful? When channel generates errors frequently

Pros

- more efficient capacity use
 - only retransmit missing packets

Cons

- receiver buffering
 - to store out-of-order packets
- more complicated buffering & protocol processing
 - to keep track of missing out-of-order packets

Tradeoff again between buffering/processing complexity and capacity

Sequence numbers HOW USED IN PRACTICE

Sequence #s in practice

How large must seq # space be?

- depends on window size

Example

- $\text{ seq } \# \text{ space} = [0, 2^4 1]$
- window size = 8

Window

Acks sent

Acks not received, times out and retransmits seq #0-7

Receiver willing to accept seq #0-7 Sender sending seq# 0-7 but different packets!

Solution: seq # space must be large enough to cover both sender + receiver windows. I.e., >= 2x window size

Sequence #s in practice

What are they counting?

- bytes, not packets
 - sending packets but counting bytes
 - so seq #s do not increase incrementally

Sequence # space

- finite
 - e.g., 32 bits so 0 to 2³²-1 values
 - must wrap around to 0 when hit max seq #
- TCP initial seq # is randomly chosen from space of values
 - security (harder to spoof)
 - to prevent confusing segments from different connections
 - different operating systems set differently: can fingerprint machines

TCP OVERVIEW

Transmission Control Protocol (TCP)

RFCs: 793,1122,1323, 2018, 2581

Main transport protocol used in Internet, provides

- mux/dmux: which packets go where
- connection-oriented, point-to-point
 - 2 hosts set up connection before exchanging data, tear down after
 - bidirectional data flow (full duplex)
- flow control: don't overwhelm receiver
- congestion control: don't overwhelm network
- reliable: resends lost packets, checks for and corrects errors
- in-order: buffers data until sequential chunk to pass up
- byte stream: no msg boundaries, data treated as stream



How does TCP provide these services?

Using many techniques we already talked about

Sliding window

- congestion and flow control determine window size
- seq #s are byte offsets

Cumulative ACKs but does not drop out-of-order packets

- only one retransmission timer
 - intuitively, associate with oldest unACKed packet
- timeout period
 - estimated from observations
- fast retransmit
 - 3 duplicate ACKs trigger early retransmit

TCP is not perfect but works pretty well!

TCP segment structure



No	. Time	Source	Destination			
	42 4.878920	172.217.11.10	<pre>vmanfredismbp2.wireless.wesleyan.edu</pre>			
	44 4.879137	outlook-namnortheast2.offi	<pre>vmanfredismbp2.wireless.wesleyan.edu</pre>			
	46 4.879346	<pre>vmanfredismbp2.wireless.we</pre>	outlook-namnortheast2.office365.com			
	Internet Protocol	Version 4, Src: outlook-namno	rtheast2.office365.com (40.97.120.226), Dst: v			
	Transmission Contro	ol Protocol, Src Port: 443 (4	43), Dst Port: 52232 (52232), Seq: 0, Ack: 1,			
	Source Port: 443	3				
	Destination Port	: 52232				
	[Stream index: 0)]				
	[TCP Segment Len	n: 0]				
	Sequence number:	<pre>0 (relative sequence numbries)</pre>	per)			
	Acknowledgment n	number: 1 (relative ack num	nber)			
	Header Length: 3	32 bytes				
1	Flags: 0x012 (SY	N, ACK)				
	000	<pre>. = Reserved: Not set</pre>				
	0	<pre>. = Nonce: Not set</pre>				
	0	<pre>. = Congestion Window Reduced</pre>	(CWR): Not set			
	0	<pre>. = ECN-Echo: Not set</pre>				
	0	<pre>. = Urgent: Not set</pre>				
	1	<pre>. = Acknowledgment: Set</pre>				
	0 = Push: Not set					
	0 = Reset: Not set					
		. = Syn: Set				
		0 = Fin: Not set				
	[TCP Flags: *	*****A**S*]				
	Window size valu	ie: 8190				
	[Calculated wind	low size: 8190]				
1	Checksum: 0xcb80	[validation disabled]				
	Urgent pointer:					
	▶ Uptions: (12 byt	es), Maximum segment size, No	D-Operation (NOP), window scale, No-Operation			
	► ISFN/A(K analvei					
00	00 78 4f 43 73 43	26 3c 8a b0 1e 18 01 08 00 4	15 20 x0CsC& <e< th=""></e<>			
00	00 34 32 41 40	00 eb 06 7e eb 28 61 78 e2 8	31 85 .42A@ ~.(ax			
00	bb ae 01 bb cc	08 ay a2 40 d9 59 5a 86 d8 8 00 02 04 05 50 01 02 02 04 0	30 12 M.YZ			
00		00 02 04 05 50 01 05 03 04 0				
100			••			

34

-

TCP SEQ #S AND ACK #S

TCP seq. numbers, ACKs

Sequence #s

 byte stream # of first byte in segment's data

Acknowledgements

- seq # of next byte
 expected from other side
- cumulative ACK
- Q: how does receiver handle out-of-order segments?
 - TCP spec doesn't say
 - up to implementer

Outgoing segment from sender



TCP ACKs

Cumulative ACKs (but different than in Go-Back-N)

- ACKs what receiver expects next, not last packet received
 - implicitly also ACKs everything up to sequence number received
- only 1 retransmission timer (for first pkt in window)
 - sender retransmits only first pkt in window if no ack when timer expires

Sequence #s are not sequential: counting bytes not packets



TCP seq. numbers, ACKs

Sequence numbers are synchronized during connection set-up



Host 1



segment from receiver?

Host 2

ansmission Control Protocol, Src
Source Port: 443 Convention: SVN
Destination Port: 54573
[Stream index: 2] and Fin take I
[TCP Segment Len: 0] Dyte of Seq #
Sequence number: 3712814908 Space
Acknowledgment number: 59452066
Header Length: 40 bytes
Flags: 0x012 (SYN, ACK)
Window size value: 14480

Transmission Control Protocol, Src Pc Source Port: 443 Destination Port: 54573 [Stream index: 2] [TCP Segment Len: 0] Sequence number: 3712814909 Acknowledgment number: 59452278 Header Lengthr 32 bytes Flags: 0x010 (ACK) Window size value: 122 [Calculated window size: 15616] [Window size scaling factor: 128]

Segment size

Max length of IP packet in bytes

- MTU: Maximum Transmission Unit
- 1500 bytes if Ethernet used as link layer protocol

Max length of TCP data in bytes

- MSS: Maximum Segment Size
- MSS = MTU IP hdr TCP hdr
 - TCP header >= 20bytes



TCP segment sent when either it is full (meets MSS) or not full but timeout occurs