

Transport Layer: TCP and UDP

Raj Jain

Washington University in Saint Louis

Saint Louis, MO 63130

Jain@wustl.edu

Audio/Video recordings of this lecture are available on-line at:

<http://www.cse.wustl.edu/~jain/cse473-22/>

Student Questions

Note: Questions discussed during regular Q&A sessions are marked with □. New questions asked during Exam 1 review are marked with ❖



- ❑ Transport Layer Design Issues:
 - Multiplexing/Demultiplexing
 - Reliable Data Transfer
 - Flow control
 - Congestion control
- ❑ UDP
- ❑ TCP
 - Header format, connection management, checksum
 - Congestion Control
- ❑ **Note:** This class lecture is based on Chapter 3 of the textbook (Kurose and Ross) and the figures provided by the authors.

Student Questions

- ❖ I am still a bit confused about acknowledgement number/ACK in general. When is the ACK field seqnum +1, and when is it equal to seqnum?

It depends on the transport protocols. TCP designers decided to make “Ack n” mean “I have received n-1st byte and I am waiting for nth byte.” In most other protocols “Ack n” means “I have received nth packet and I am waiting for n+1st packet”.



Transport Layer Design Issues

1. Transport Layer Functions
2. Multiplexing and Demultiplexing
3. Error Detection: Checksum
4. Flow Control
5. Efficiency Principle
6. Error Control: Retransmissions

Student Questions

- What is the difference between SDUs and PDUs?

---- From Slide 3-7 ----

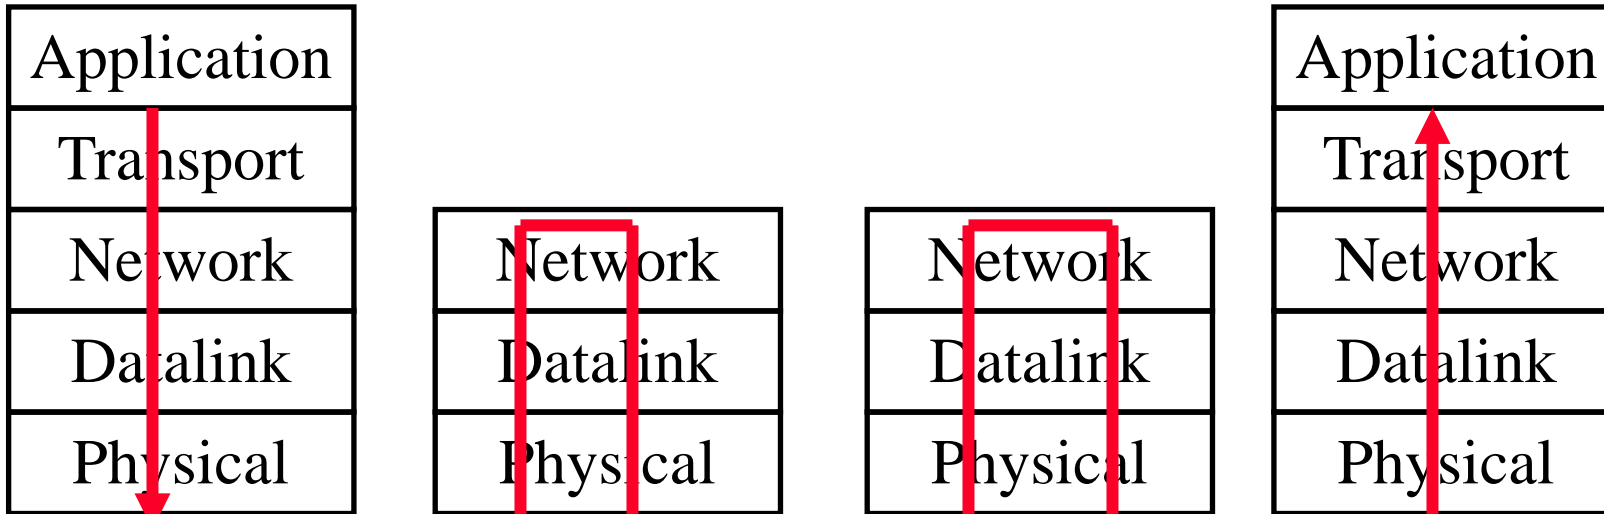
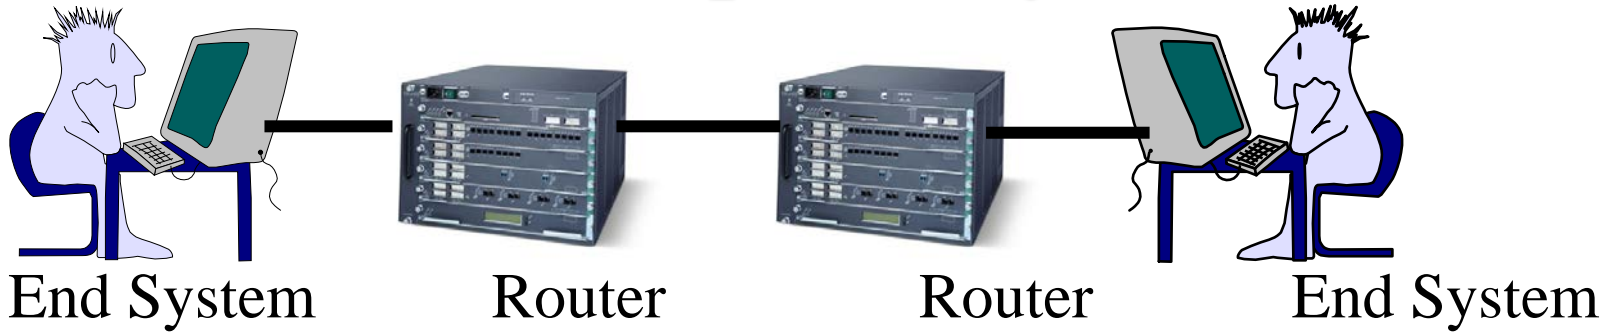
- Is the multiplexing of ports usually done by frequency or time?

Multiplexing is on the link not at the port. Packets from different applications use different port numbers to avoid confusion at the destination. After receiving, the packets are serialized in time, processed by lower layers, and queued by the port number for the application to process.

- Is it possible for an application to use more than one port? If so, does it give them any benefit in terms of data transfer rates?

One application could be talking to many applications, each requiring a different port. For example, a movie player may talk to video server, caption server, quiz server, and audio server. All 4 using different ports. It can also talk to two video servers – both using the same port but at different destinations.

Transport Layer

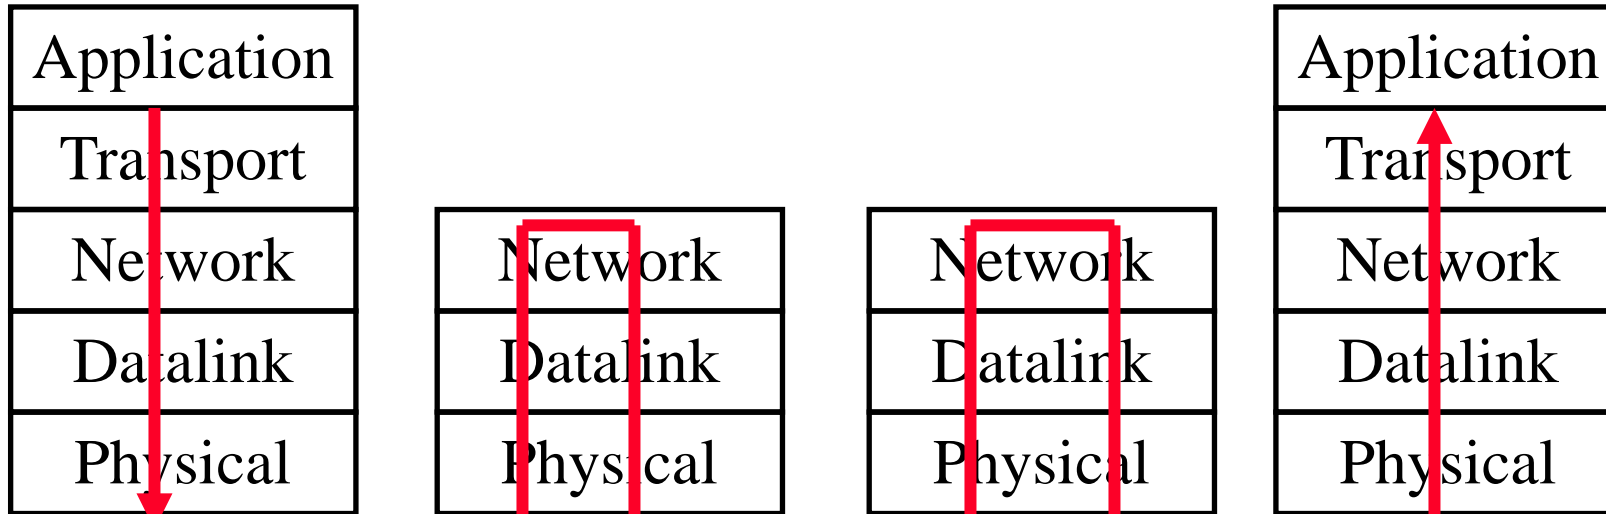
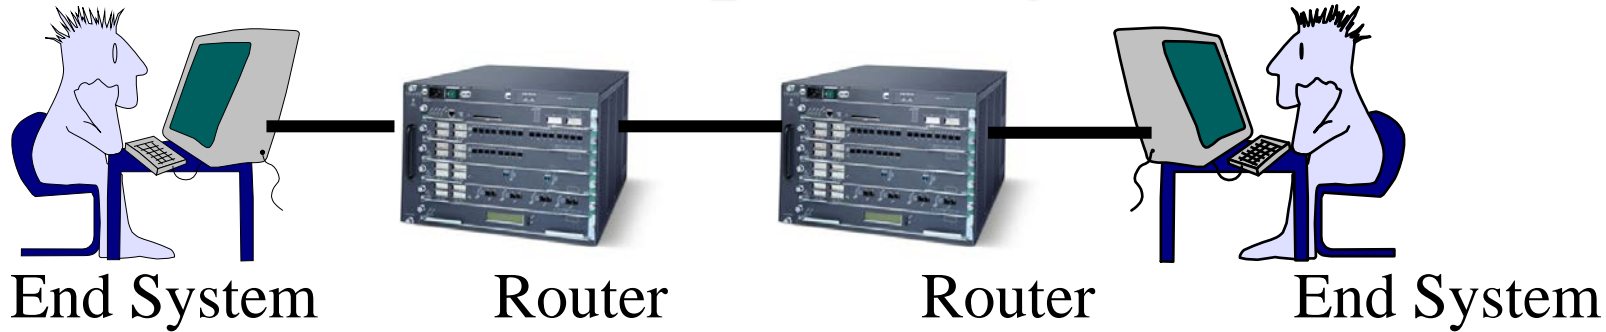


- ❑ Transport = End-to-End Services
Services required at source and destination systems
Not required on intermediate hops

Student Questions

- ❑ What do buffers refer to?
Incoming and outgoing packets are stored in the memory. That part of the memory is called buffer.
- ❑ Does this mean that routers cannot do Flow control & Loss detection & Congestion Control since it doesn't have a Transport layer?
Technically yes. But it can react to its own buffer overflow.
- ❑ What do hops mean?
1 Hop = 1 Link
- ❑ What does the program running on the router look like? Is it a Linux-based program?
Proprietary OS. Could be Linux. There is no standard.
- ❑ What hardware does Network Interface Controller (NIC) run on?
- ❑ NIC is the network card connected to the CPU. For example, USB dongles for Ethernet are NICs.
- ❑ Is port 80 here a socket port? Is a socket just address plus port?
Please see Chapter 2 Q&A video and slides.

Transport Layer



- ❑ Transport = End-to-End Services
Services required at source and destination systems
Not required on intermediate hops

Student Questions

- ❖ When a packet goes through a router, does the router check the layer-4 header and layer-5 header of the packet? Since layer-4 header and layer-5 header do not contain IP addresses, is the routing process done in layer-3 by router?

Yes, the router looks only at the layer-3 header. layer 4-5 headers are part of the data field and are not interpreted by the network layer. The network layer is responsible for getting the packet to the destination IP address. Therefore, routers are responsible for routing.

Transport Layer Functions

1. **Multiplexing and demultiplexing:** Among applications and processes at end systems
2. **Error detection:** Bit errors
3. **Loss detection:** Lost packets due to buffer overflow at intermediate systems (Sequence numbers and acks)
4. **Error/loss recovery:** Retransmissions
5. **Flow control:** Ensuring destination has buffers
6. **Congestion Control:** Ensuring network has capacity

Not all transports provide all functions

Student Questions

□ What kinds of transports don't provide functions?
There are many transports. Some do not provide some functions. Each uses a different method for each function. This is a comprehensive but not complete list of functions provided by various transports.

□ How does the computer know if bits or packets were lost if it doesn't have the original file to compare it to? Is it using the checksum?

A packet is received but there is no buffer to store it.

□ Can the intermediate nodes have any capability restrictions with regards to retransmitting data, or is it all the transport protocol?

No restrictions. Most wireless routers do retransmit. More about it in Chapter 6 on Datalink.

□ What transport does not provide one or more of the above functions?

UDP. See Slide 2-10.

□ Is there a function that is mandatory for all transports?

No. Continued on Slide 3-6

Protocol Layers

□ Top-Down approach

Application	HTTP	FTP	SMTP	P2P	DNS	Skype
Transport	TCP				UDP	
Internetwork	IP					
Host to Network	Ethernet	Point-to-Point			Wi-Fi	
Physical	Coax	Fiber	Wireless			

Student Questions

---From Slide 3-5---

- Is this a comprehensive list of all possible functions?

No. This list is a good start.

- Are networks considered "better" if they implement more functions?

No. More functions ⇒ More Delay ⇒ More Cost

- Does congestion control include flow control?

No. But they are similar.

-----For Slide 3-6-----

- Can an application send out a request using TCP (e.g., HTTP request) on one port number and request that the response be sent back to a different port number?

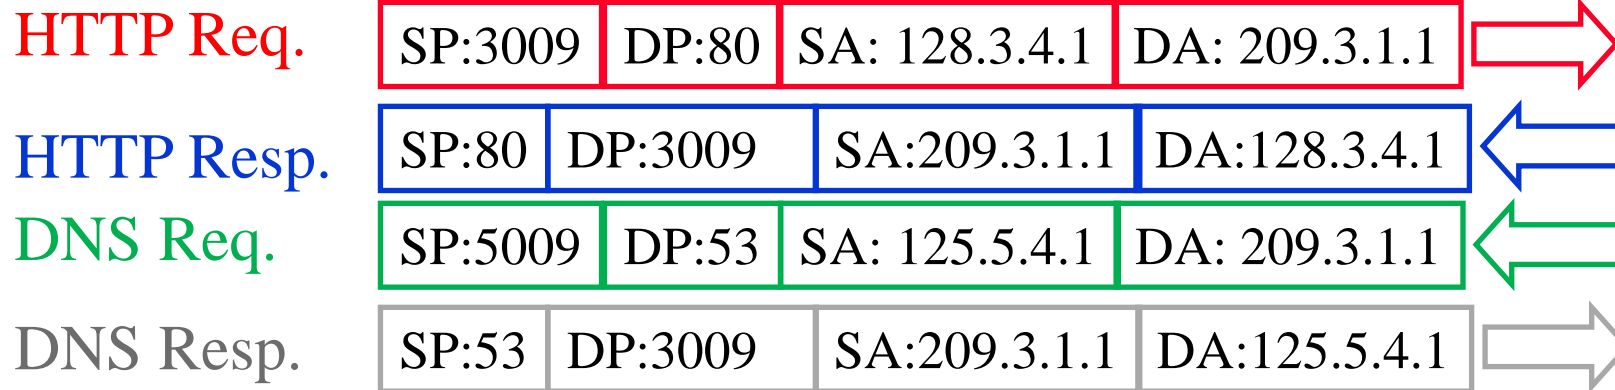
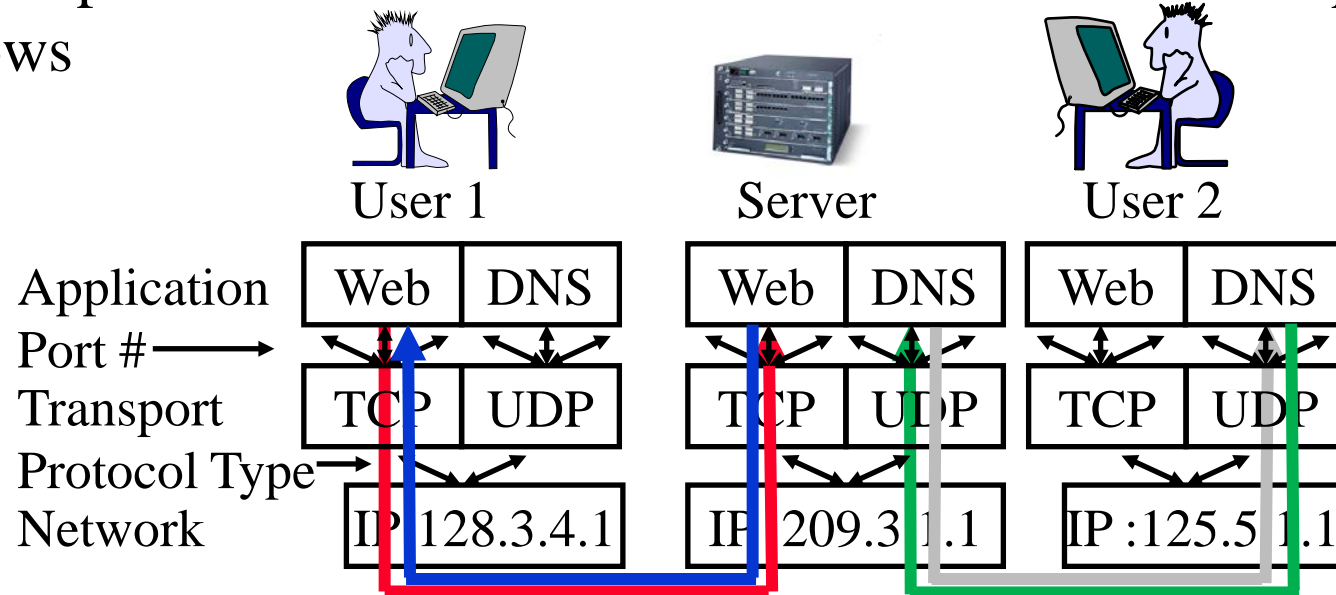
The source port numbers are randomly picked from a range outside the standard port #.

- If the sender does not specify a port number, which port does the transport layer resolve that message to?

The packet is dropped. Application s/w at the source puts in the default port # if needed.

Multiplexing and Demultiplexing

- Transport **Ports** and Network **addresses** are used to separate flows



Ref: http://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers

Student Questions

- Can you explain this diagram again? Can't see where the laser was pointing to from the video. *Sure.*
- Can applications send and receive over the same port simultaneously? *Yes.*

- We often use the default port number to communicate with the server. On a TCP or UDP-based server, wouldn't it cause any problems if many people were talking to the same port number?

Port numbers are like doors. Many people can arrive through the same door.

- Can a service port accept multiple requests within the same time frame? How is this traffic handled and manipulated by the port?

Multiple packets can enter the same port.

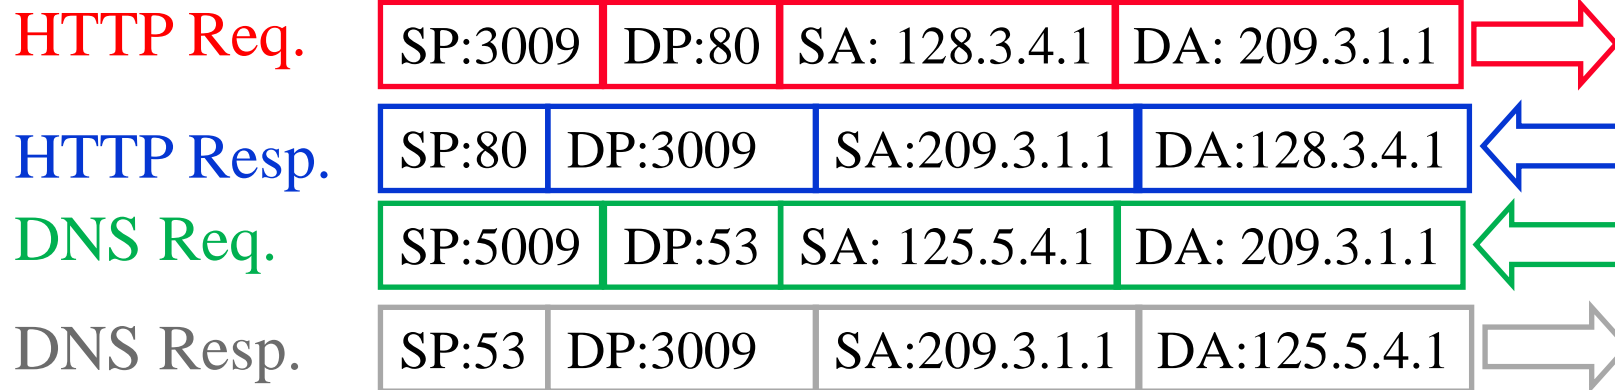
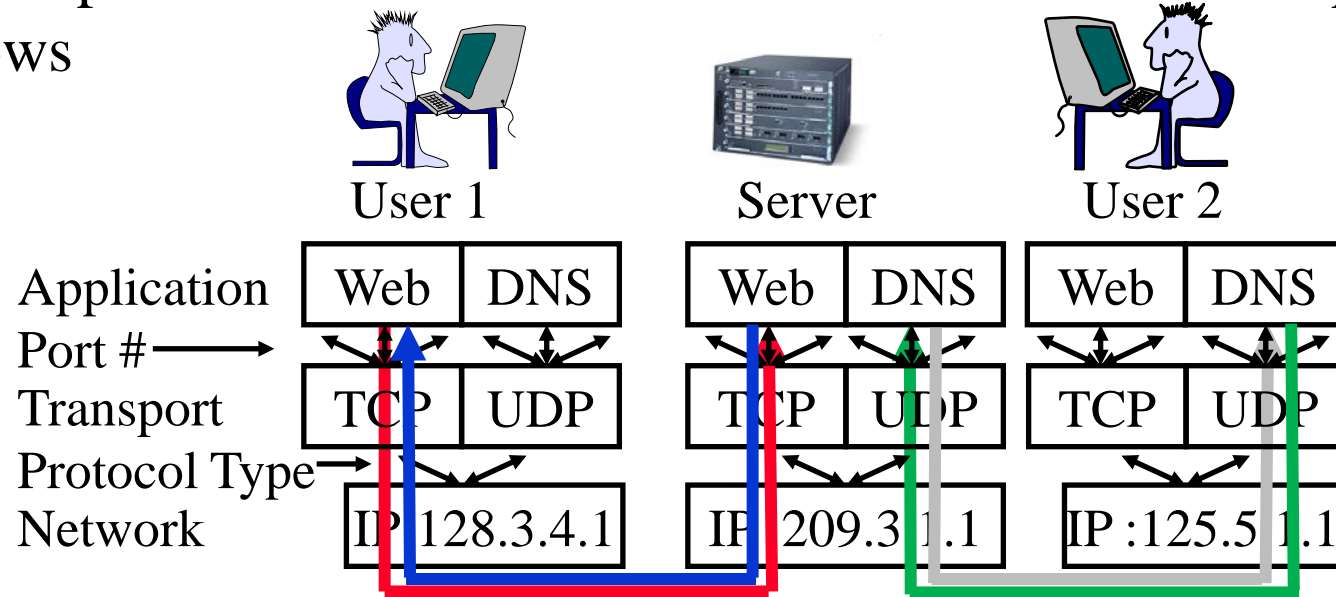
- Is the transport port the same as the application port?

Yes. Transport port is the Service Access Point for transports

---- Continued on Slide 3-3 ----

Multiplexing and Demultiplexing

- Transport **Ports** and Network **addresses** are used to separate flows



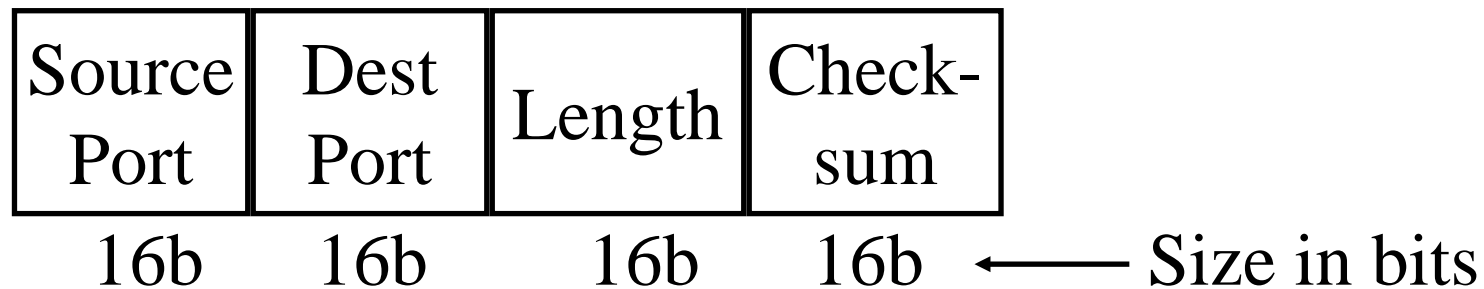
Ref: http://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers

Student Questions

- ❖ Can we cover slide 3-7 since I am having a hard time understanding the diagram?
- ❑ *Sure.*

User Datagram Protocol (UDP)

- ❑ Connectionless end-to-end service
- ❑ Provides multiplexing via ports
- ❑ Error detection (Checksum) optional. Applies to **pseudo-header** (same as TCP) and UDP segment. If not used, it is set to zero.
- ❑ No error recovery (no acks). No retransmissions.
- ❑ Used by network management, DNS, Streamed multimedia (Applications that are loss tolerant, delay sensitive, or have their own reliability mechanisms)



Student Questions

- ❑ You mention RFC. what is that?

RFC = Request for Comments

RFCs started out as memos that were used by initial Internet research team to inform other team members about their design decisions. These are numbered and are still used. Some of these are "informational" while others are "standards" that specify a protocol. TCP is specified in RFC 793:

<https://tools.ietf.org/html/rfc793>

- ❑ Why use UDP over TCP even if you don't require reliability?

Reliability costs in terms of delay. In overloaded networks, packets may be lost causing receiver to wait for retransmissions. Sometimes it is better to move on and not wait for retransmissions. This is the case with audio/video.

- ❑ Why would you use checksum with UDP since there is no way to recover the data if there is an error?

To throw away corrupted information.

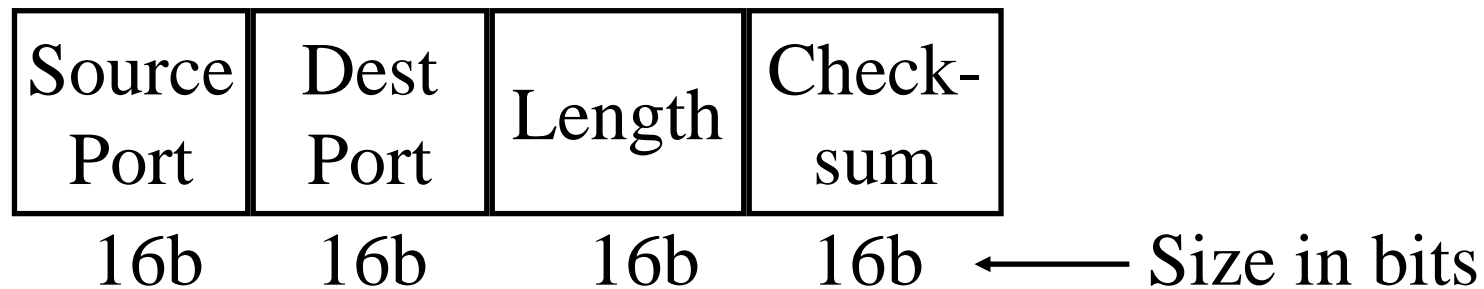
- ❑ Can UDP be based on circuit switching? Circuit switching is required a complete connection in physics, but UDP is a connectionless end-to-end service. What is the difference between "connection" here?

Connection = Set up before talking

-----Continued on Slide 3-15 -----

User Datagram Protocol (UDP)

- ❑ Connectionless end-to-end service
- ❑ Provides multiplexing via ports
- ❑ Error detection (Checksum) optional. Applies to **pseudo-header** (same as TCP) and UDP segment. If not used, it is set to zero.
- ❑ No error recovery (no acks). No retransmissions.
- ❑ Used by network management, DNS, Streamed multimedia (Applications that are loss tolerant, delay sensitive, or have their own reliability mechanisms)

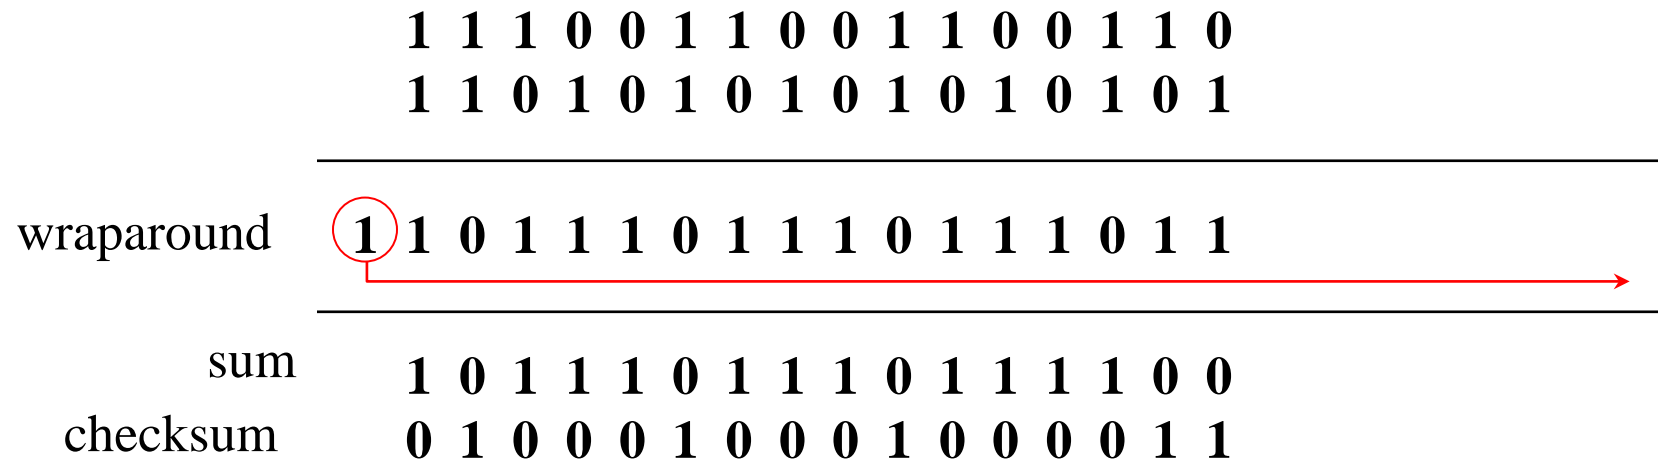


Student Questions

- ❖ Do most UDP transmissions opt to use the checksum or not?
Yes. It is almost universal. No apps can work with bad data.
- ❖ When do we take pseudo header to the checksum?
Always for UPD and TCP.

Error Detection: Checksum

- ❑ **Cyclic Redundancy Check (CRC):** Powerful but generally requires hardware
- ❑ **Checksum:** Weak but easily done in software
 - **Example:** 1's complement of 1's complement sum of 16-bit words with overflow wrapped around



At receiver the sum is all 1's and the checksum is zero.

Student Questions

- ❑ What is an error that can arise from getting -0
In all cases -0 is considered the same as 0. So there is no error.
- ❑ Is 1's complement better than 2's?
No. But, some tricks work with one representation but not the other and so they are used in different context.

-
- ❑ Since in 2's complement, the smallest number (the largest negative number) has a greater magnitude than the most positive, on fixed-sized ALU's, how do you negate that number since there is no positive equivalent?

We don't use that number.

- ❑ Besides checksum, any other ways for CRC?

Many other ways. More in Chapter 6.

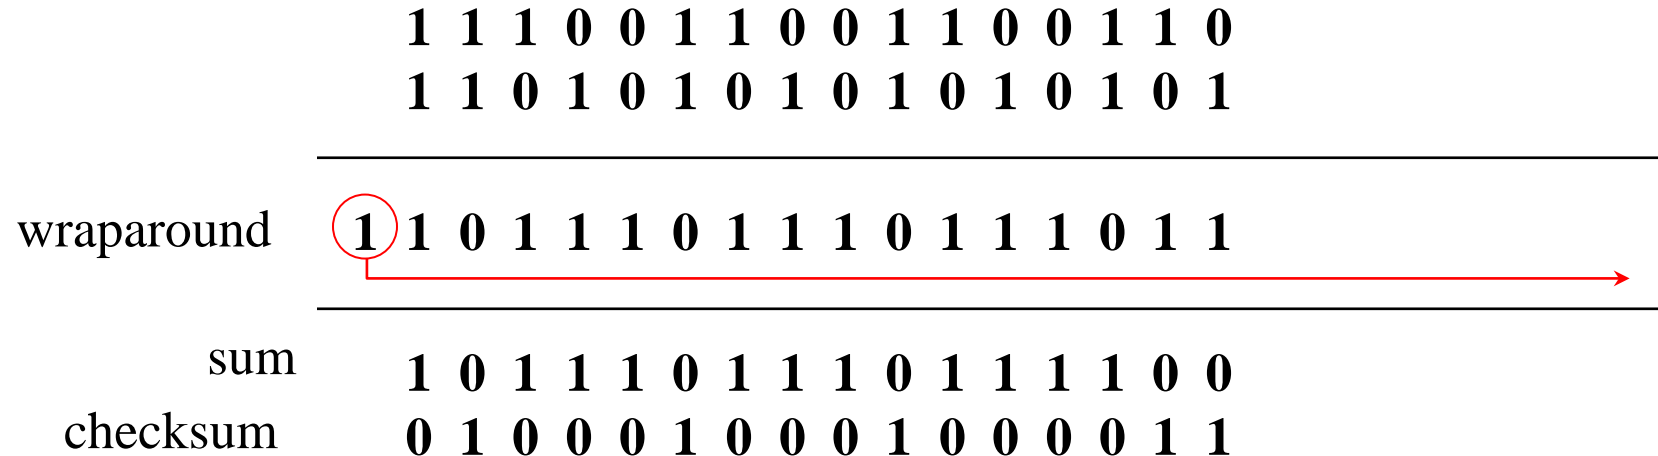
- ❑ Are the two 16-bit words predefined by the application? Trying to understand the purpose checksum provides.

This is just an example. All bits in the packets are arranged as rows of 16-bits. Then the checksum is computed and added to the header.

-----Continued on Slide 3-12 -----

Error Detection: Checksum

- ❑ **Cyclic Redundancy Check (CRC):** Powerful but generally requires hardware
- ❑ **Checksum:** Weak but easily done in software
 - **Example:** 1's complement of 1's complement sum of 16-bit words with overflow wrapped around



At receiver the sum is all 1's and the checksum is zero.

Student Questions

- ❖ Why is it that we take the 1's complement sum of the header and message for checksum instead of just checking if the two checksums are equal?

At the receiver the checksum comes out zero. Rather than comparing if it is 25, you can add -25 to the message so that the receiver will get 0.

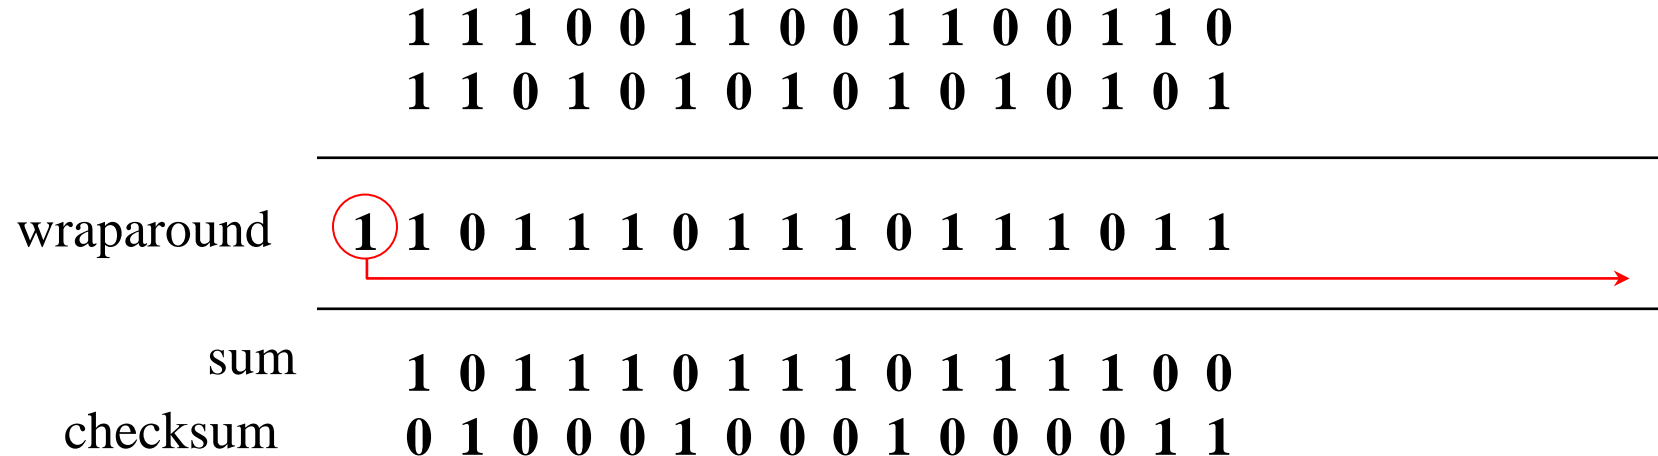
- ❖ if the checksum(1's complement) becomes wrong when sent to the receiver, will it still be wrong even if other words are all correct?

Yes. We do know which byte is in error, the entire message will be considered in error.

- ❑ Why do we need 1's complement of checksum instead of just add data up?
- ❑ *See the first question above.*

Error Detection: Checksum

- ❑ **Cyclic Redundancy Check (CRC):** Powerful but generally requires hardware
- ❑ **Checksum:** Weak but easily done in software
 - **Example:** 1's complement of 1's complement sum of 16-bit words with overflow wrapped around



At receiver the sum is all 1's and the checksum is zero.

Student Questions

- ❑ Can you please explain, what does "1's complement of 1's complement sum" on slide 3-9 means? Isn't 1's complement of 1's complement sum, just the sum itself?

No. All operations here are done using "1's complement arithmetic" Please do not confuse the name of the method with the operation.

Checksum = Complement(Sum)

You are thinking

Checksum =

Complement(Complement(Sum))

Error Detection: Checksum

- ❑ **Cyclic Redundancy Check (CRC):** Powerful but generally requires hardware
- ❑ **Checksum:** Weak but easily done in software
 - **Example:** 1's complement of 1's complement sum of 16-bit words with overflow wrapped around



At receiver the sum is all 1's and the checksum is zero.

Student Questions

- ❖ Book pp. 199 - 3.3.2 UDP checksum. This question applies to TCP and UDP checksums. If we have two bytes: 0000 0000 and 0000 1111 and they are corrupted such that the new (corrupted) bytes are 0000 0010 and 0000 1101, won't they still pass the checksum? The probability of this is quite low but it could happen. Does TCP have any recourse for this?

This is the case of a two-bit error. Two-bit errors may or may not be detected by the TCP checksum. If not detected, the message will be considered correct and delivered to the application unless other inconsistencies prevent it from delivery.

- ❖ Can we go through an example of calculating TCP checksum?

Sure.

1's Complement

2's Complement: -ve of a number is complement+1

- ❑ 1 = 0001 -1 = 1111
- ❑ 2 = 0010 -2 = 1110
- ❑ 0 = 0000 -0 = 0000

1's complement: -ve of a number is it's complement

- ❑ 1 = 0001 -1 = 1110
- ❑ 2 = 0010 -2 = 1101
- ❑ 0 = 0000 -0 = 1111

2's Complement sum: Add with carry. Drop the final carry, if any.

$$6-7 = 0110 + (-0111) = 0110 + 1001 = 1111 \Rightarrow -1$$

1's complement sum: Add with carry. Add end-around carry back to sum

- ❑ $6-7 = 0110 + (-0111) = 0110+1000 = 1110 \Rightarrow -1$

Complement of 1's complement sum: 0001

Checksum: At the transmitter: 0110 1000, append 0001

At the receiver: 0110 1000 0001 compute checksum of the full packet
= complement of sum = complement of 1111 = 0000

Ref: https://en.wikipedia.org/wiki/Ones%27_complement

Washington University in St. Louis

<http://www.cse.wustl.edu/~jain/cse473-22/>

©2022 Raj Jain

Student Questions

- ❑ For the scope of this exam, when would we need to use 2's complement?

2's complement are presented here to contrast with 1's complement. We may use 2's complement during CRC discussion in Chapter 6.

- ❑ If we are sending 16-bit packets and we want to use a checksum, how many bits of the packet are the actual data, and how many bits of the packet is the complement sum?

16-bit packets would be too small for anything. A packet consists of many 16-bit words. One extra 16-bit word is added as a checksum.

- ❑ Does 2's complement related to sign bit in this case?

No. Two's complement relates to the entire word not just the sign bit.

- ❑ For the example of checksum, how can we get the appended number(0001 in this case) at the transmitter?

By computing the complement of the 2's complement sum of all words in the packet.

-----Continued on Slide 3-11 -----

Homework 3A: Checksum

[6 points] Consider the following two 16-bit words: ABCD 1234

- A. What is the checksum as computed by the sender
- B. Add your answer of Part A to the end of the packet and show how the receiver will compute the checksum of the received three 16-bit words and confirm that there are no errors.
- C. Now assume that the first bit of the packet is flipped due to an error. Repeat Part B at the receiver. Is the error detected?

Student Questions

- For the encoding of ABCD, would it just be decimal values 10-14 in binary? Assuming 0-9 are the same in their binary encoding.

$$9_{16} = 9_{10} = 1001_2$$

$$A_{16} = 1010_2$$

$$19_{10} = 13_{16} = 00010011_2$$

- Is the first bit of the packet referring to the entire packet, including headers? Or first bit of a message?

We have not used "message" in this homework. The packet consists of two words indicated plus one word of checksum. Total three words.

- In Homework3A, why do we use 1's Complement? Why not 2's Complement?
The homework does not specify which complement to use. Please use the complement that will be used by UDP/TCP.

-----From Slide 3-10 -----

- 1's complement addition almost seems pointless since it's so easy to convert two binary numbers to base 10, normally add, then convert back to 1's complement. Is there any reason to avoid doing it this way?

Yes. Computers do not know decimal arithmetic. They only know binary arithmetic.

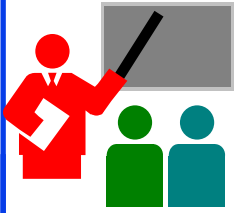
Homework 3A: Checksum

[6 points] Consider the following two 16-bit words: ABCD 1234

- A. What is the checksum as computed by the sender
- B. Add your answer of Part A to the end of the packet and show how the receiver will compute the checksum of the received three 16-bit words and confirm that there are no errors.
- C. Now assume that the first bit of the packet is flipped due to an error. Repeat Part B at the receiver. Is the error detected?

Student Questions

- ❖ Book Page 211: On the receiver side, how to compute the checksum of the acknowledgment packet being sent to layer 3?
- ❑ *As usual, the ack packet will be assembled and checksum computed at the transmitter and checked at the receiver. Nothing special about ack packets. Most acks come with data as data packets anyway.*



UDP: Summary

1. UDP provides flow multiplexing using port #s
2. UDP optionally provides error detection using the checksum
3. UDP does not have error or loss recovery mechanism

Read Sections 3.1, 3.2, and 3.3, Do R1-R8, P2-P5

Washington University in St. Louis

<http://www.cse.wustl.edu/~jain/cse473-22/>

©2022 Raj Jain

Student Questions

-----From Slide 3-9 -----

- Does UDP throw the entire packet? Or only the part has errors?

We don't know which part has error. The entire packet is dropped.

- If something is wrong with both the data bits and checksum bits, is that possible that the wrong checksum could just match the wrong data, and we will not detect the error in data?

There is a small probability of undetected errors.

- Why is UDP used instead of TCP?
Because there is no waiting for lost packet recovery.

- Why is this called "pseudo-header" and not just "header"?

It is prepared in the memory but not added to the packet.

- Is the purpose of a pseudo-header only to carry the checksum? *Yes*
- Does TCP also use a pseudo-header, or is it only used by UDP? *Both*
- Can an application use TCP and UDP at the same time? *Yes*

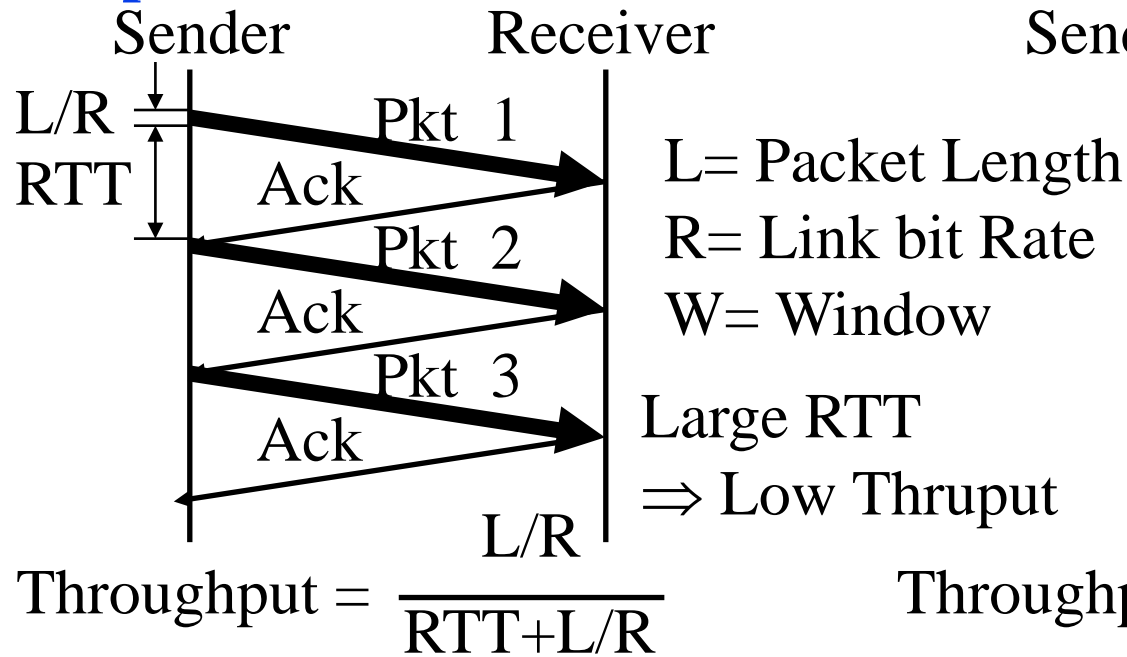


Flow Control

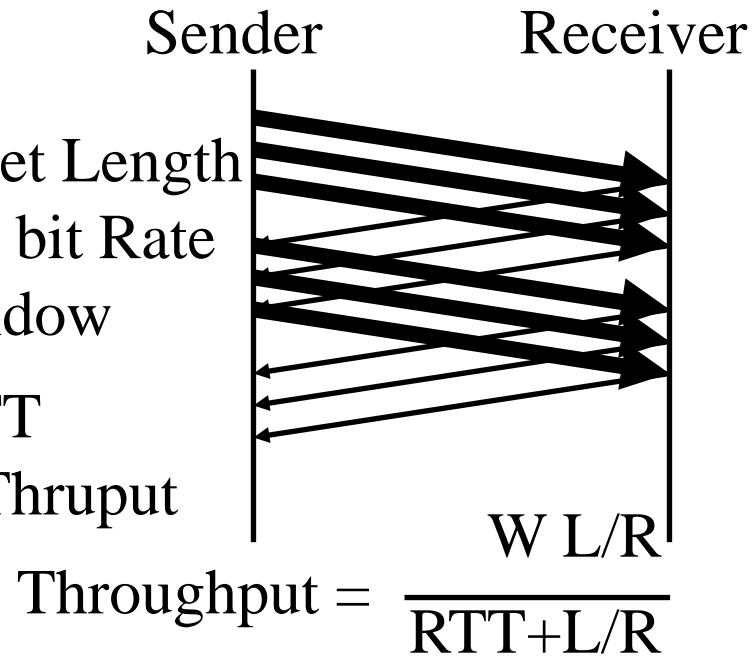
Flow Control Goals:

1. Sender does not flood the receiver,
2. Maximize throughput

Stop and Wait Flow Control



Window Flow Control

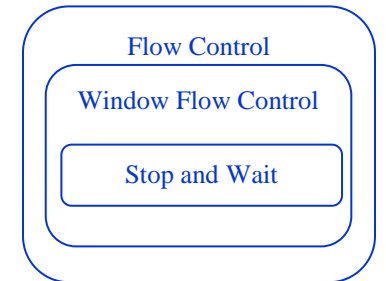


Ref: Textbook Section 3.4.2

Student Questions

- What's the actual difference between Stop and Wait and Window Flow Control

The following Venn diagram explains that Stop and Wait is one of the of flow control methods.



- Why, in this case $(L/R)/(RTT=L/R)$ is the throughput? I mean, according to slide 15 is also called utilization.

Utilization is between 0 and 1.

Throughput is in bits/second.

- The unit of the throughput discussed before is bps. Why does the throughput become a ratio now? *Ratio=% throughput*

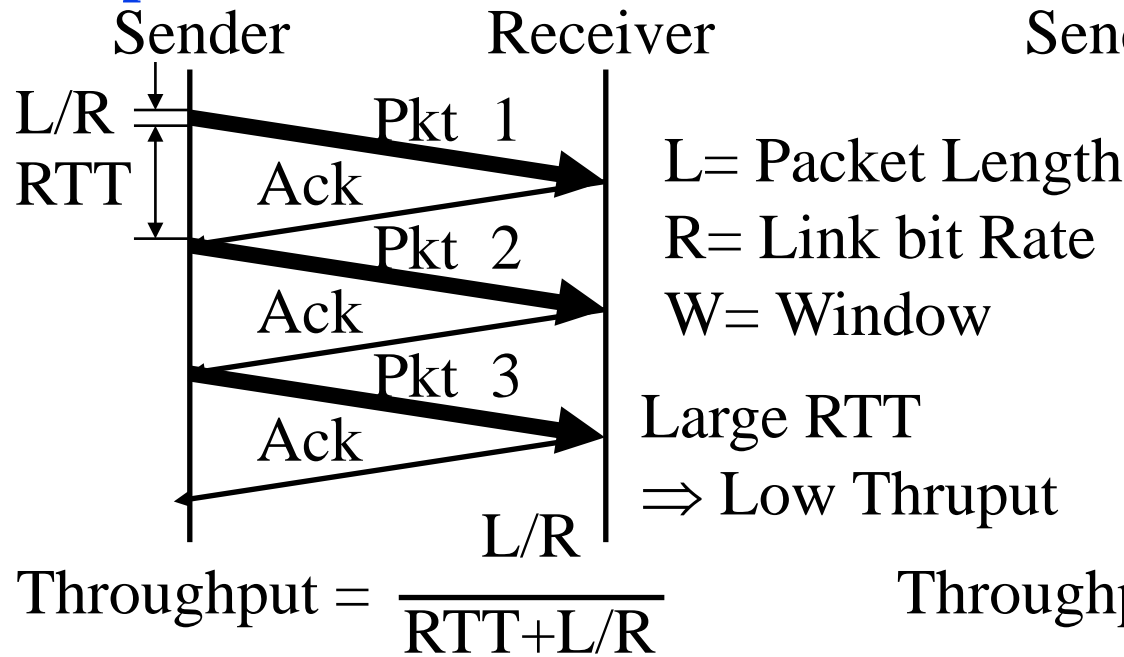


Flow Control

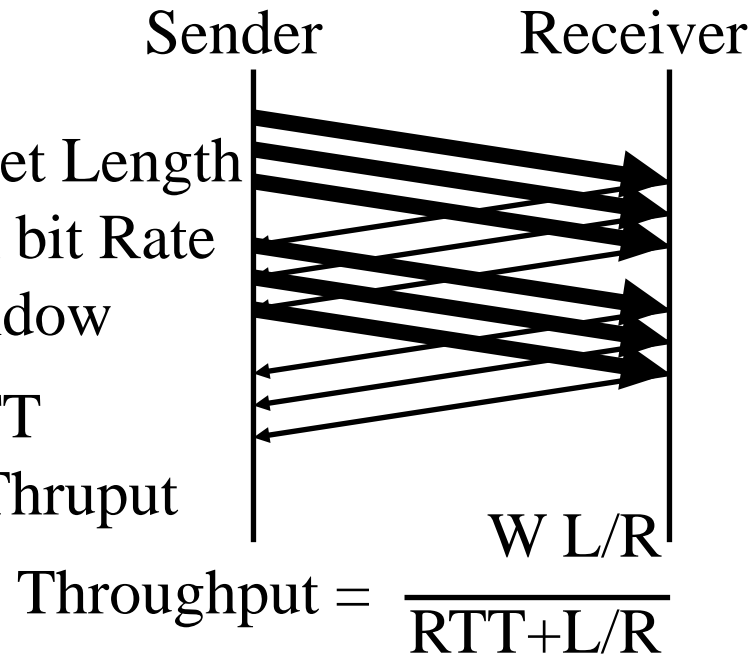
Flow Control Goals:

1. Sender does not flood the receiver,
2. Maximize throughput

Stop and Wait Flow Control



Window Flow Control

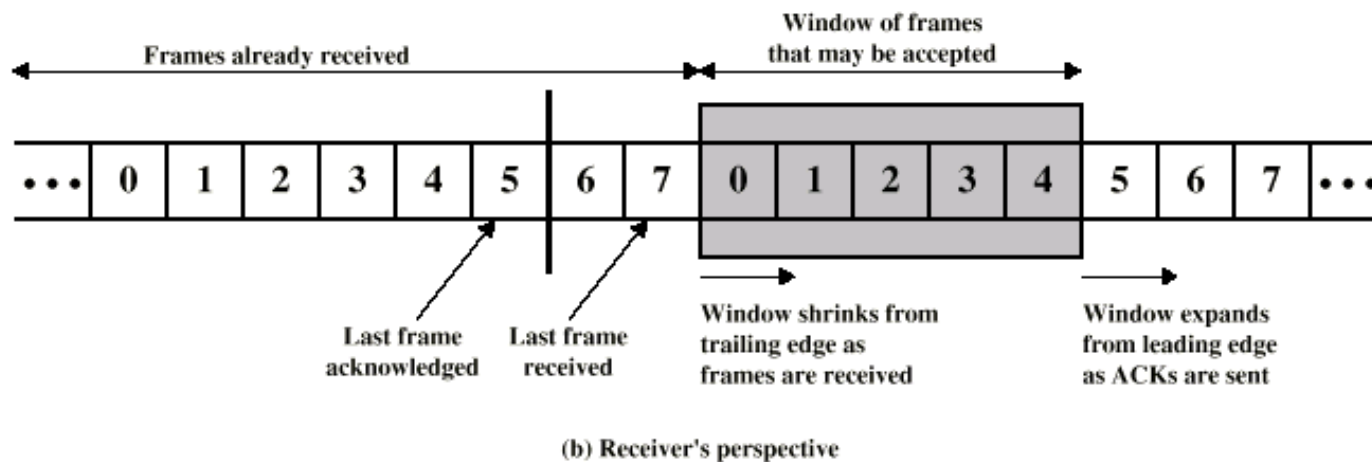
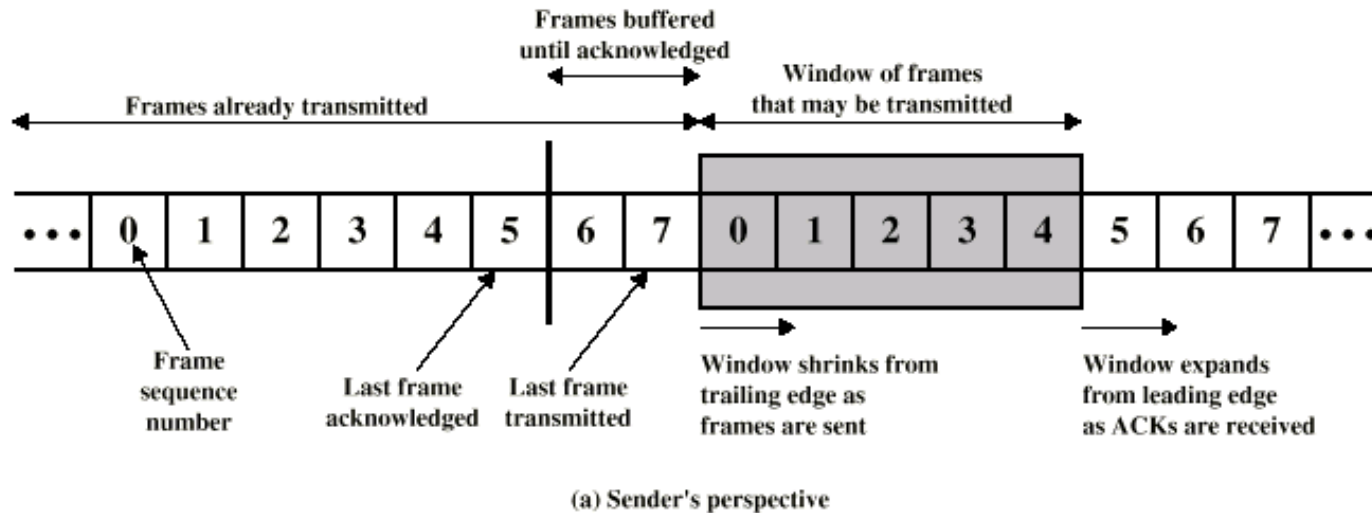


Student Questions

- ❖ In the textbook, the GBN and SR protocol is attributed to the reliable transmission service. But the class attributed them to the flow control technology. Does flow control have to use techniques for the reliable data transfer service?

Any violation of flow control will result in a packet loss which is the same as unreliability. Bit errors also result in unreliability but will not violate SR or GBN. So flow control is a more correct classification for SR or GBN. This is just a necessary condition for reliable transmissions but not sufficient.

Sliding Window Diagram



Student Questions

- Is there a maximum size of the window, if 6 and 7 get acked, will the sender's window expand by 2 frames, or wait till it sends 2 frames before expanding.

Window shows the packets that can be sent but have not been sent. The maximum size is determined by the receiver. The acks generally include instructions for how many more packets the receiver can receive. This slide shows a special case in which each ack that is acking n packets allows n more packets. This is not necessary.

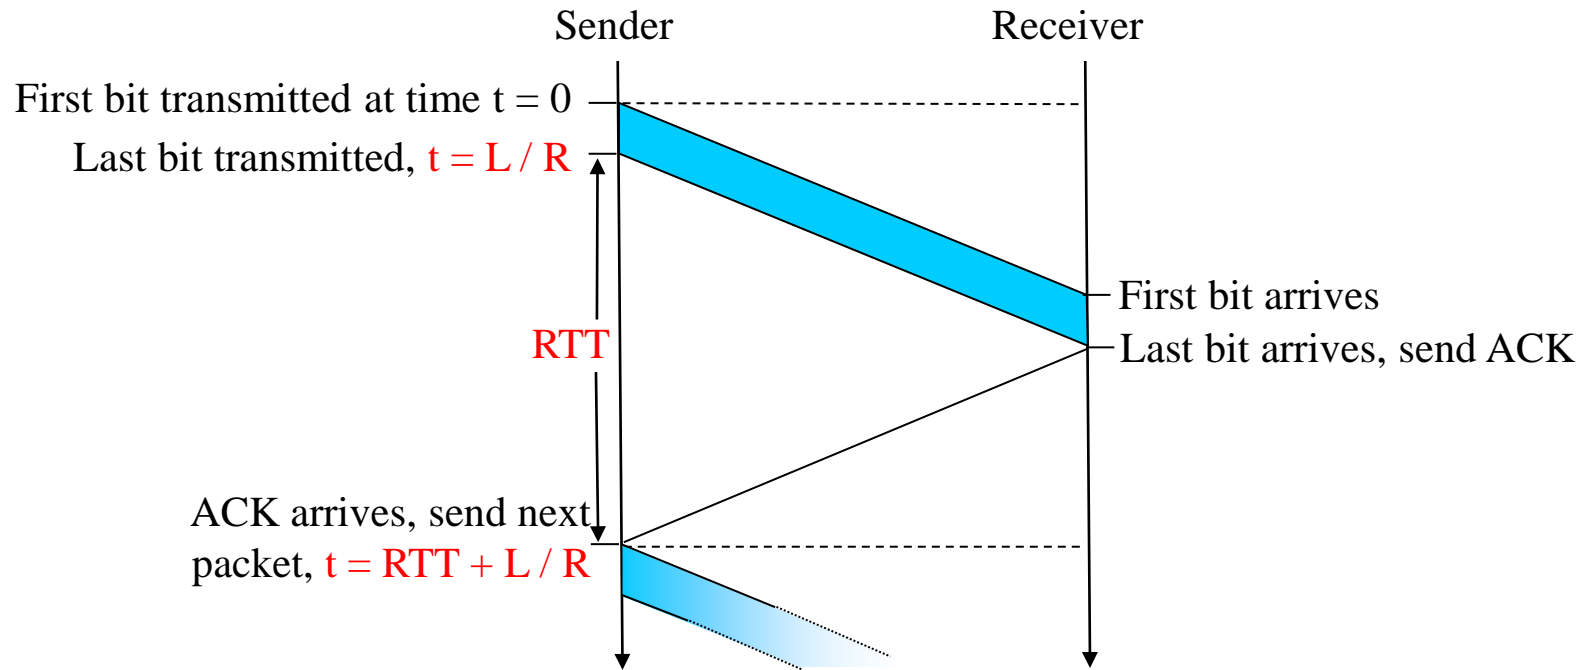
- Could you explain the sliding window diagram again since we can't see the pointer on the video?

Sure.

- Is the size of the window fixed? Is the number of unacknowledged packets limited by a certain threshold?

Yes. That is called "window size."

Stop and Wait Flow Control



$$\text{Utilization } U = \frac{L/R}{RTT + L/R} = \frac{t_{\text{frame}}}{2t_{\text{prop}} + t_{\text{frame}}} = \frac{1}{2\alpha + 1}$$

Here, $\alpha = t_{\text{prop}}/t_{\text{frame}}$

Student Questions

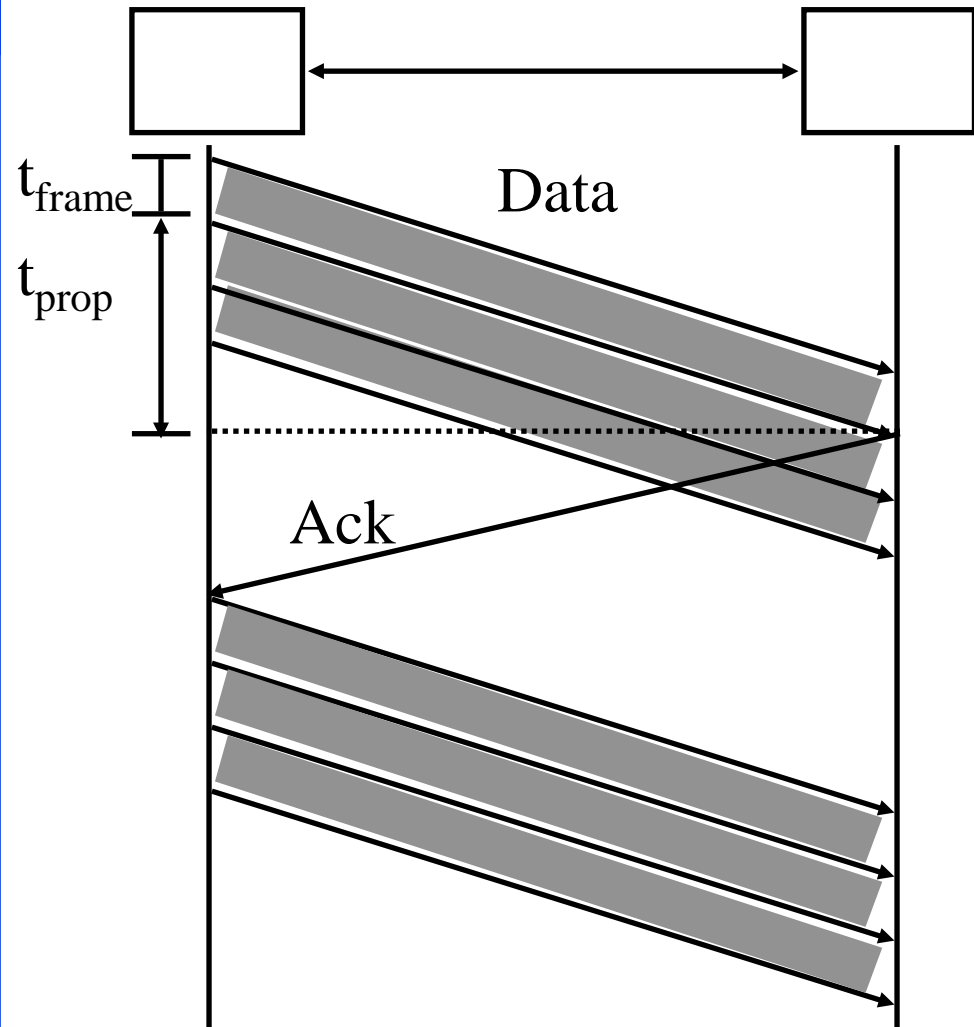
- Why is the round trip time equal to $2 \cdot t_{\text{prop}}$? Is this including processing and queuing delays?
 t_{prop} here is actually one-way delay it includes propagation, processing and queueing.

---- From Slide 3-8 ----

- Those streaming multimedia applications are often considered not to be data loss sensitive. They may use UDP in the transport layer with respect to the textbook, but why do many of them are actually using TCP instead?

The application may have more than Video such as file size (which is data). TCP is used for data.

Sliding Window Protocol Efficiency



$$U = \frac{W t_{\text{frame}}}{2t_{\text{prop}} + t_{\text{frame}}}$$

$$= \begin{cases} \frac{W}{2\alpha + 1} \\ 1 \text{ if } W > 2\alpha + 1 \end{cases}$$

Here, $\alpha = t_{\text{prop}}/t_{\text{frame}}$

$W=1 \Rightarrow$ Stop and Wait

Student Questions

Utilization: Examples

Satellite Link: One-way Propagation Delay = 270 ms

RTT=540 ms

Frame Size $L = 500$ Bytes = 4 kb

Data rate $R = 56$ kbps $\Rightarrow t_{\text{frame}} = L/R = 4\text{kb}/56\text{kbps} = 0.071\text{s} = 71$ ms

$\alpha = t_{\text{prop}}/t_{\text{frame}} = 270/71 = 3.8$

$U = 1/(2\alpha+1) = 0.12$

❑ Short Link: 1 km = 5 μs (Assuming Fiber 200 m/ μs),

Rate=10 Mbps,

Frame=500 bytes $\Rightarrow t_{\text{frame}} = 4\text{k}/10\text{M} = 400$ μs

$\alpha = t_{\text{prop}}/t_{\text{frame}} = 5/400 = 0.012 \Rightarrow U = 1/(2\alpha+1) = 0.98$

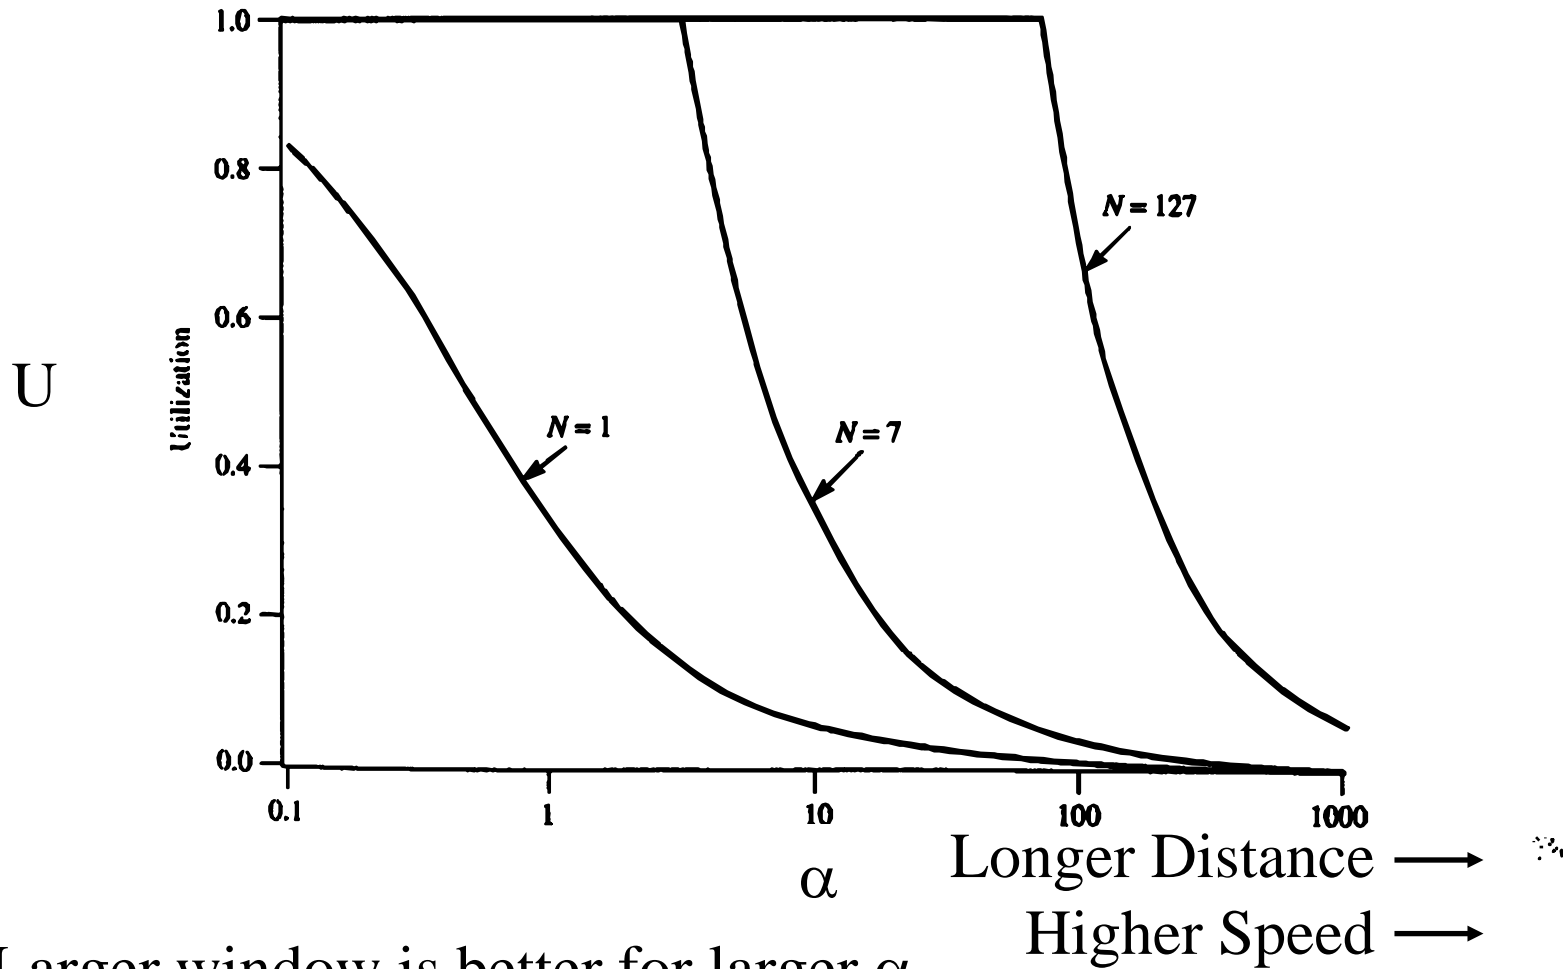
Note: The textbook uses RTT in place of t_{prop} and L/R for t_{frame}

Student Questions

❖ When calculating U should we use RTT or one-way propagation delay?

α uses one-way propagation delay

Effect of Window Size



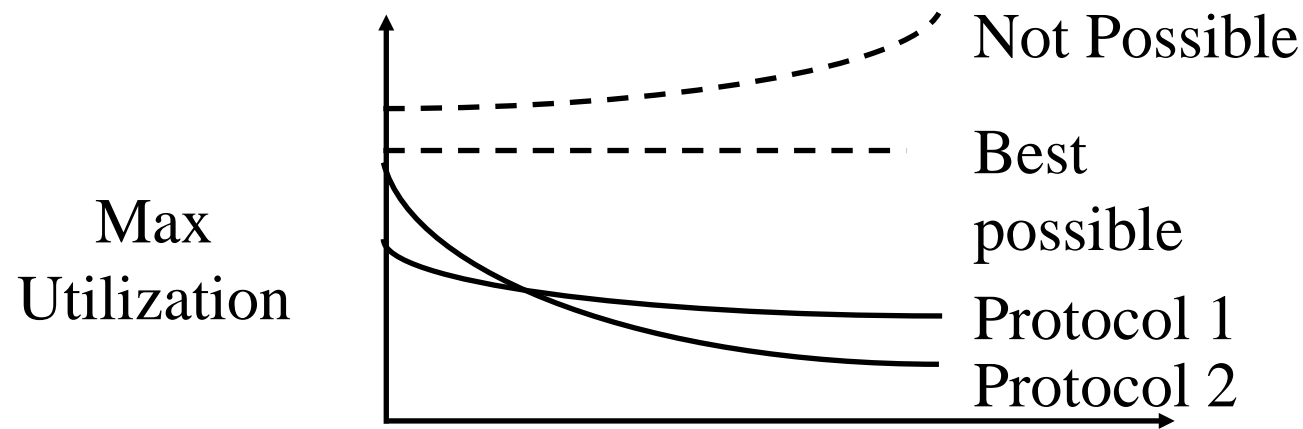
- Larger window is better for larger α

Student Questions

- Why is a larger window better for a larger α ?
The graph shows that the larger N gives better utilization for larger α .
Large α
 \Rightarrow Long distance network
For full utilization, you must fill up the entire path with packets so you will need more packets as the distance increases. If the path is empty, the throughput is low. If the path is overfilled, delay increases due to queueing.

Efficiency Principle

- For **all** protocols, the maximum utilization (efficiency) is a *non-increasing* function of α .



$$\alpha = \frac{t_{\text{prop}}}{t_{\text{frame}}} = \frac{\text{Distance/Speed of Signal}}{\text{Bits Transmitted /Bit rate}}$$

$$= \frac{\text{Distance} \times \text{Bit rate}}{\text{Bits Transmitted} \times \text{Speed of Signal}}$$

Student Questions

- Why do you use the term "non-increasing" when describing the shape of the graph? Why not "constant"? Is there a specific reason?

Non-increasing = Decreasing + Constant

- Is efficiency the same as utilization?

They are different but equal if 100% of resources are used.

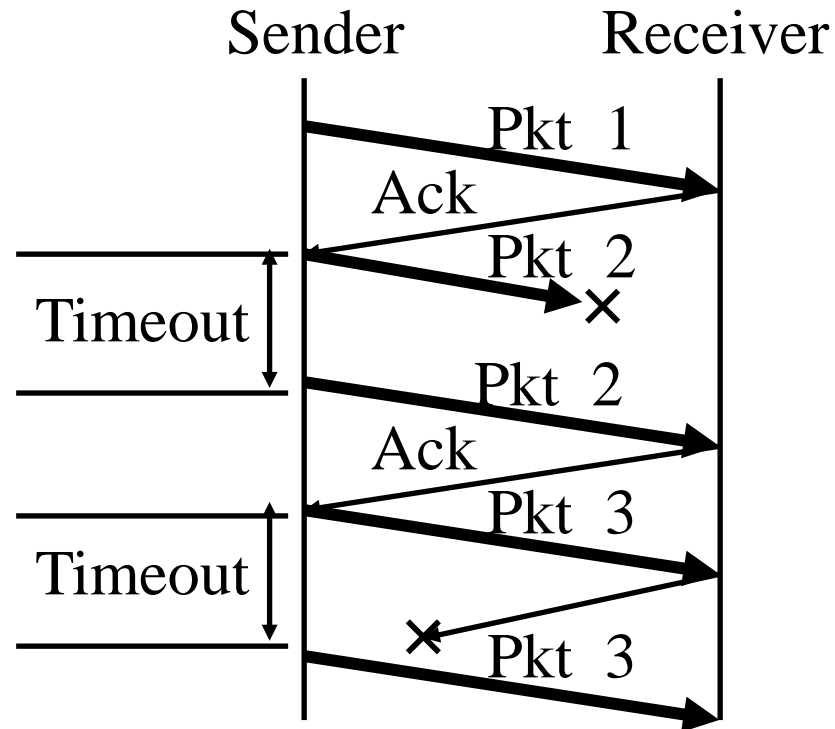
$$\text{Utilization} = \frac{\text{Output traffic}}{\text{Capacity}}$$

$$\text{Efficiency} = \frac{\text{Output traffic}}{\text{Resources used}}$$

Error Control: Retransmissions

- ❑ Error Control = Error Recovery
- ❑ Retransmit lost packets \Rightarrow Automatic Repeat reQuest (ARQ)

Stop and Wait ARQ

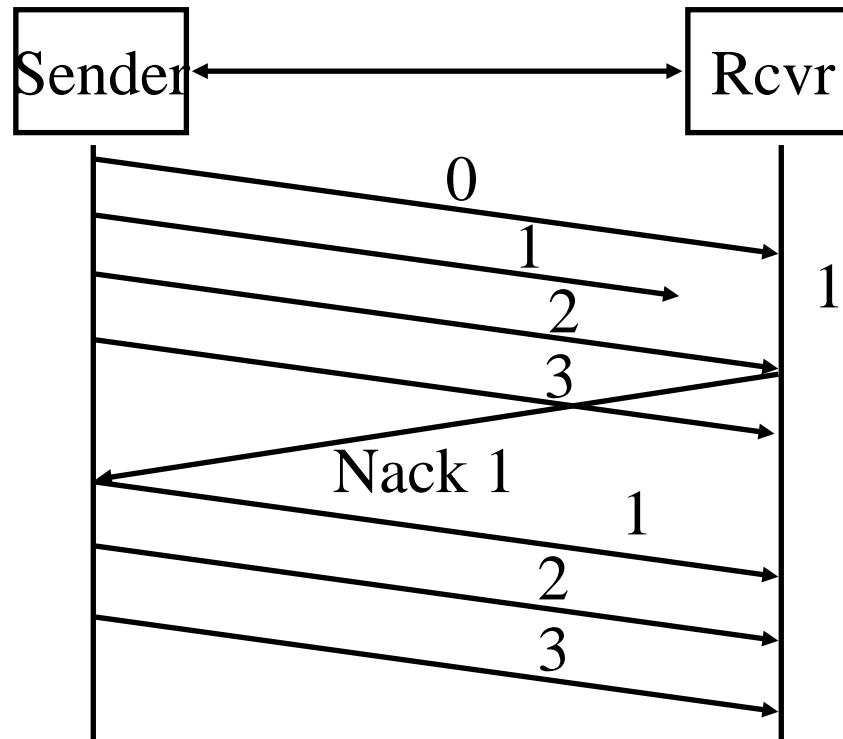


Student Questions

- ❑ Is it possible that we thought a packet was lost after a timeout, and we resend it, then we receive two Ack that this packet is received? If so, can we identify what packet is received from Ack?

This is quite common. Both packets are identical. We throw away one. But two acks are sent to ensure that the ack was not lost.

Go-Back-N ARQ



- ❑ Receiver does not cache out-of-order frames
- ❑ Sender has to *go back* and retransmit all frames after the lost frame

Student Questions

- ❑ TCP has similarities with GBN/selective repeat, but does not explicitly use these, correct?

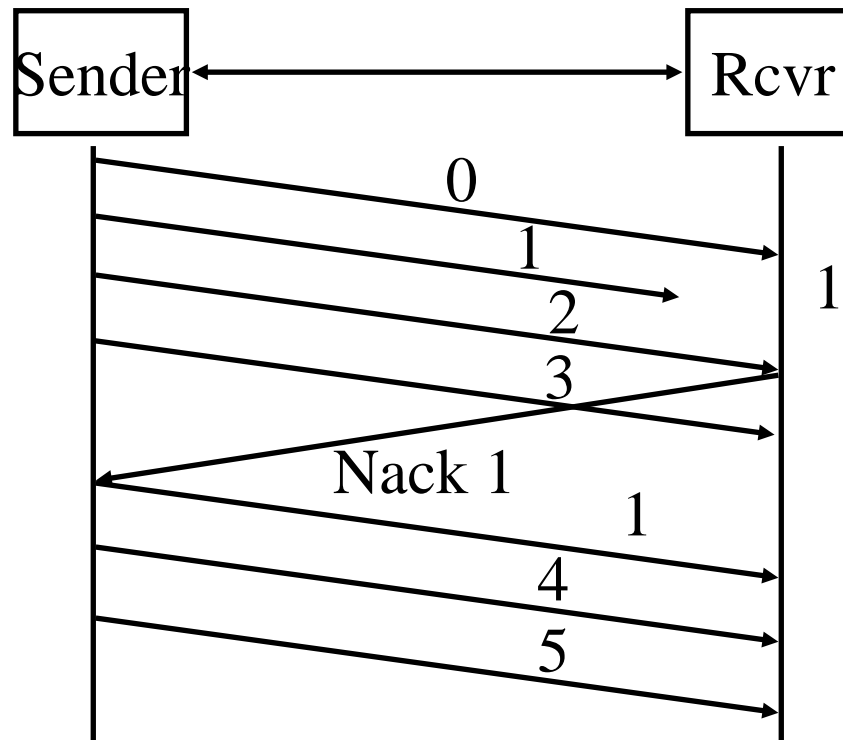
- ❑ Is there a standard on how long a message is held on to after it is sent? Could the sender respond to a Nack 0 after Message 3 is sent?

Yes. The transmitter and receiver operate asynchronously. They react to events at their end.

- ❑ How does the receiver know which segment is lost?

It keeps track of sequence numbers of received packets.

Selective Repeat ARQ



- ❑ Receiver caches out-of-order frames
- ❑ Sender retransmits only the lost frame
- ❑ Also known as selective *reject* ARQ

Student Questions

- ❑ Can you nack and ack at the same time in Selective Repeat ARQ. How would the Sender know that frame 2 has been received if the receiver doesn't send an ack on 2, and sends an ack (for 3) on 3. Or can you ack more than one frame at once?

All acks are cumulative. If you ack n, then all packets up to n have been received. Nacks are cumulative too. When you nack n, then all packets up to n-1 have been received.

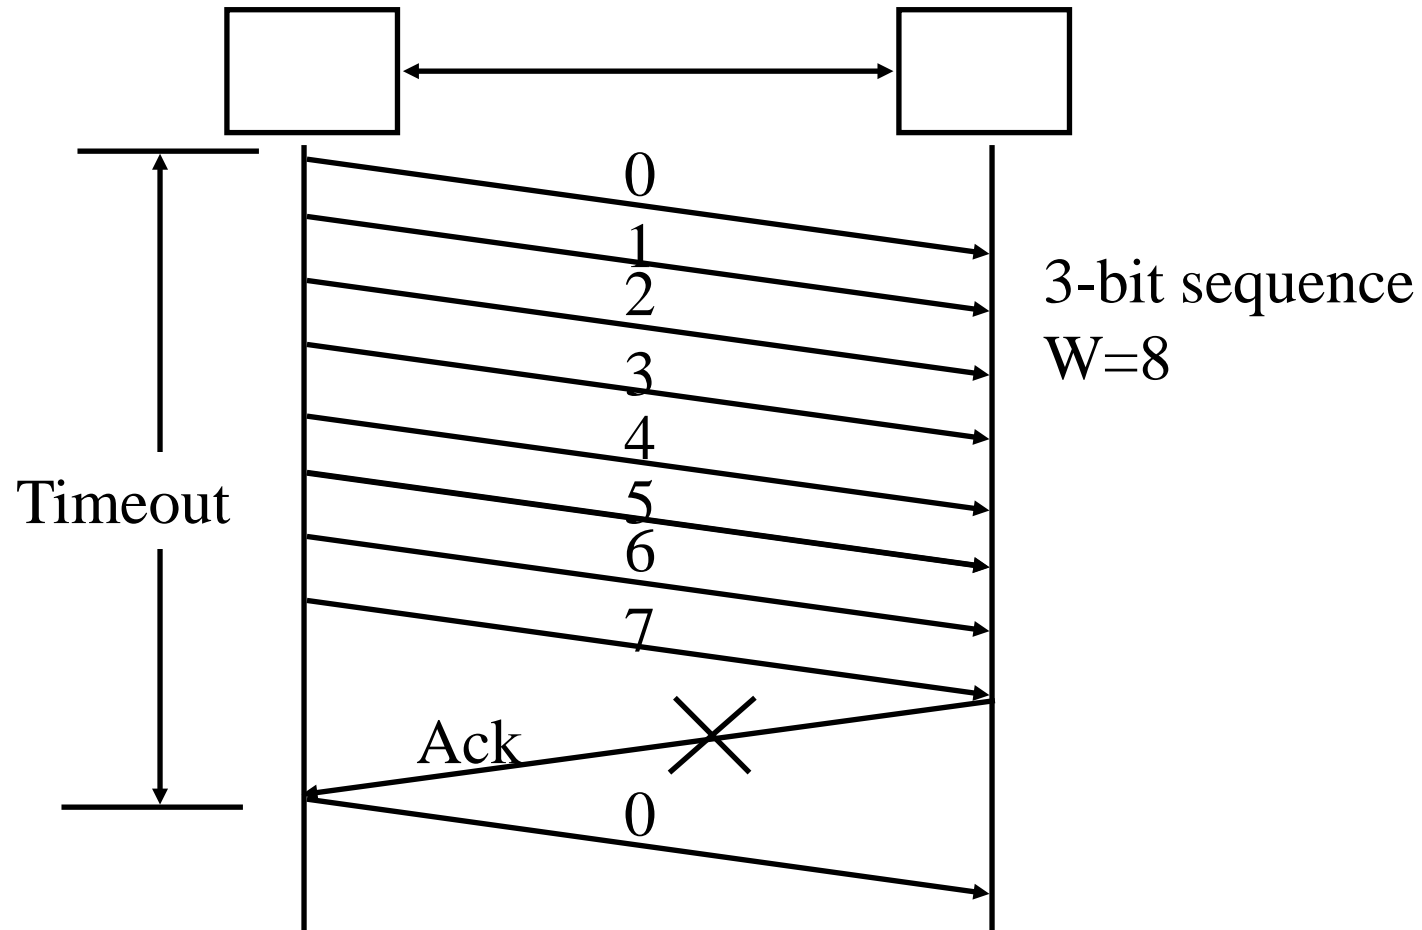
- ❑ How does the receiver know that it has missed frame 1? but don't know whether the next frame it has received is frame 2 or 3?

Every frame has a sequence number in it.

- ❖ Can selective-repeat protocol or go-back-n protocol with window size 1 used as alternating-bit protocol?

Yes.

Selective Repeat: Window Size



Sequence number space ≥ 2 window size

Window size $\leq 2^{n-1}$ with n bit sequence numbers

Student Questions

- Is an ack sent for each received packet in a window?

No. Generally, the receiver waits until it receives either sufficient number of packets or no packets are received for certain interval before sending the ack.

- How do we select the optimal window size if there are multiple possibilities?

Smallest Optimal

- In this case, does the missed packet get sent in the next window or can it be sent in any of the upcoming windows? What would be the performance implication for each of these cases?

The missed packet is sent as soon as it is realized that the packet may have been lost. The window cannot advance beyond that permitted by the packet ack'ed.

Performance: Maximum Utilization

❑ **Stop and Wait Flow Control:** $U = 1/(1+2\alpha)$

❑ **Window Flow Control:**

$$U = \begin{cases} 1 & W \geq 2\alpha + 1 \\ W/(2\alpha + 1) & W < 2\alpha + 1 \end{cases}$$

❑ **Stop and Wait ARQ:** $U = (1-P)/(1+2\alpha)$

❑ **Go-back-N ARQ:** $P = \text{Probability of Loss}$

$$U = \begin{cases} (1-P)/(1+2\alpha P) & W \geq 2\alpha + 1 \\ W(1-P)/[(2\alpha + 1)(1-P + WP)] & W < 2\alpha + 1 \end{cases}$$

❑ **Selective Repeat ARQ:**

$$U = \begin{cases} (1-P) & W \geq 2\alpha + 1 \\ W(1-P)/(2\alpha + 1) & W < 2\alpha + 1 \end{cases}$$

Student Questions

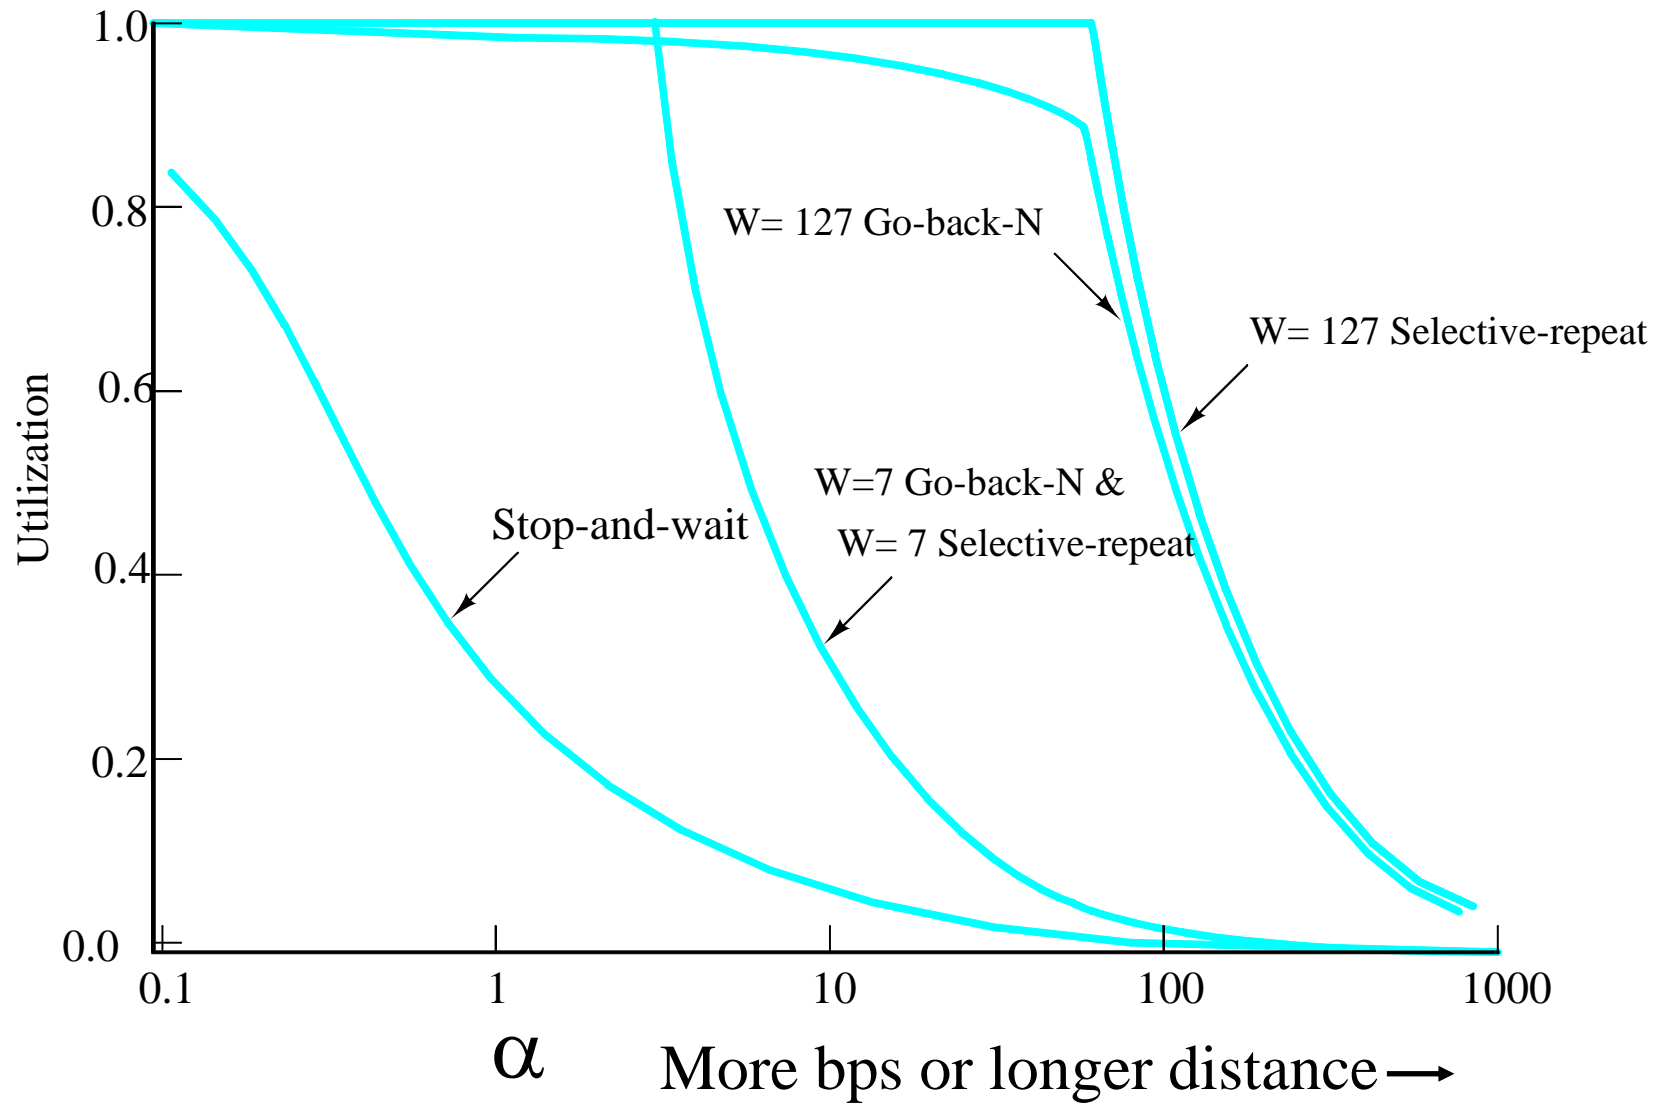
❑ How can we get the formula of go-back-N ARQ and selective repeat ARQ?

Good exercise. You can do it if you know probability theory.

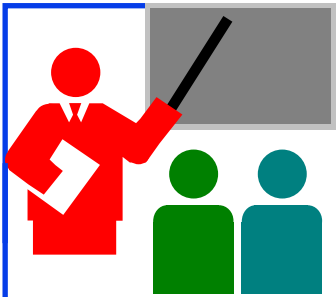
❖ Do we need to remember these formulas for the exam?

Yes

Performance Comparison



Student Questions



Transport Layer Design: Summary

1. Multiplexing/demultiplexing by a combination of source and destination IP addresses and port numbers.
2. Window flow control is better for long-distance or high-speed networks
3. Longer distance or higher speed
 \Rightarrow Larger $\alpha \Rightarrow$ Larger window is better
4. Stop and and wait flow control is ok for short distance or low-speed networks
5. Selective repeat is better than stop and wait ARQ
Only slightly better than go-back-N

Read Sections 3.4.2, 3.4.3, and 3.4.4. Do R11-13, P23, P24

Student Questions

Homework 3B: Flow Control

[8 points] Similar to problem 22 on page 292 of the textbook:

Consider the GBN protocol with a sender window size of 3 and a sequence number range of 1,024. Suppose that at time t , the next in-order packet that the receiver is expecting has a sequence number of k . Assume that the medium does not reorder messages.

Answer the following questions:

- A. What are the possible sets of sequence numbers inside the sender's window at time t ? Justify your answer.
- B. What are all possible values of the ACK field in all possible messages currently propagating back to the sender at time t ? Justify your answer.

Window Flow Control:

- C. How big window (in number of packets) is required for the channel utilization to be greater than 80% on a cross-country fiber link of 4000 km running at 20 Mbps using 2 kByte packets?

Efficiency Principle:

- D. Ethernet V1 access protocol was designed to run at 10 Mbps over 2.5 Km using 1500 byte packets. This same protocol needs to be used at 100 Mbps at the same efficiency. What distance can it cover if the frame size is not changed?

Student Questions



TCP

1. TCP Header Format, Options, Checksum
2. TCP Connection Management
3. Round Trip Time Estimation
4. Principles of Congestion Control
5. Slow Start Congestion Control

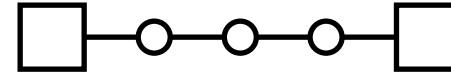
Student Questions

- ❑ UDP has no congestion control. How can it be applied in multimedia when people all use multimedia at the same time?

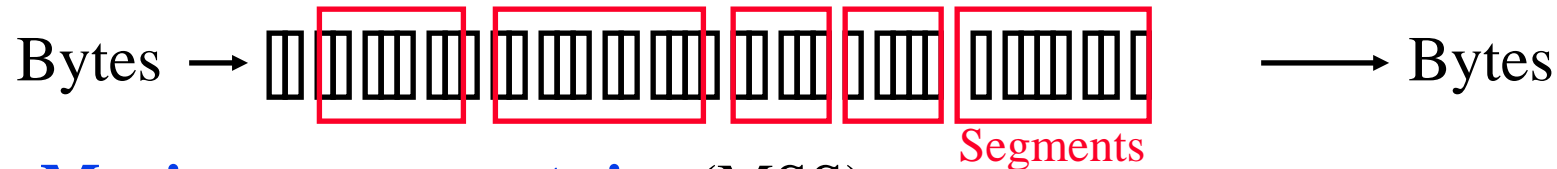
They all lose some packets. Recently congestion control has been added to UDP. But that is part of the “Recent Advances in Networking Course”

Key Features of TCP

❑ **Point-to-Point:** One sender, one receiver

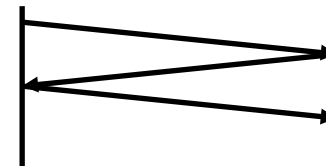


❑ **Byte Stream:** No message boundaries.
TCP makes “segments”

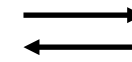


❑ **Maximum segment size (MSS)**

❑ **Connection Oriented:** Handshake to initialize states before data exchange



❑ **Full Duplex:** Bidirectional data flow in one connection

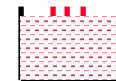


❑ **Reliable:** In-order byte delivery

❑ **Flow control:** To avoid receiver buffer overflow



❑ **Congestion control:** To avoid network router buffer overflow



Student Questions

❑ Are there any types of cyberattacks performed via overflowing a receive buffer?

Yes, buffer overflow attacks are quite common.

❑ What is the maximum segment size for TCP?

Segment=TPDU. MSS is set by the application.

Minimum MSS with IPv4 is 536B to allow sending one 512 byte block in one segment.

❑ MSS limits the size of segments, so messages will be sliced into several segments. Does the same thing happen at the link layer and internet layer?

Yes. Every layer has its own maximum PDU size.

❑ Regardless of the connection, what is the difference between the end-to-end service of UDP and the point-to-point of TCP?

Retransmission of lost PDUs.

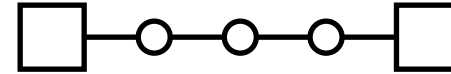
❑ Is the three-way-handshake that makes TCP a connection-oriented protocol?

Yes.

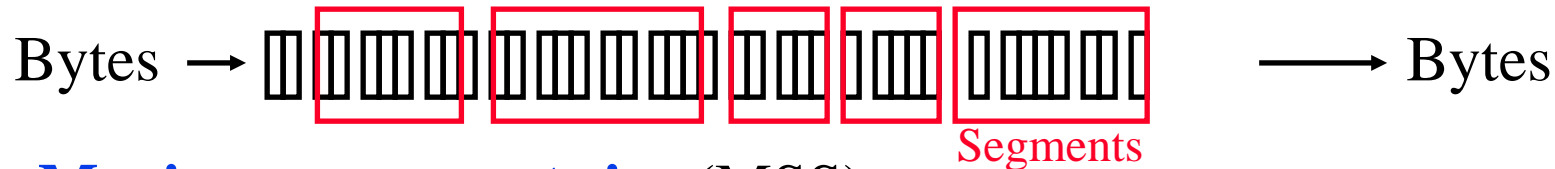
❑ What determines the length of segments? *TCP decides to send a segment either when it has enough to send or after a set period.*

Key Features of TCP

❑ **Point-to-Point:** One sender, one receiver

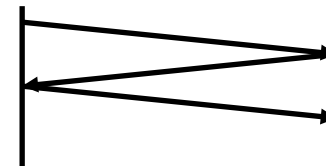


❑ **Byte Stream:** No message boundaries.
TCP makes “segments”

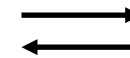


❑ **Maximum segment size (MSS)**

❑ **Connection Oriented:** Handshake to initialize states before data exchange

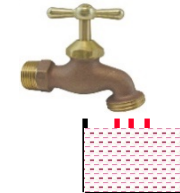


❑ **Full Duplex:** Bidirectional data flow in one connection



❑ **Reliable:** In-order byte delivery

❑ **Flow control:** To avoid receiver buffer overflow



❑ **Congestion control:** To avoid network router buffer overflow

Student Questions

❖ How is the system able to serve both packets for UDP and bytes/segments for TCP?

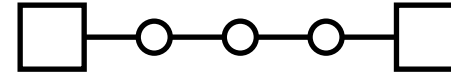
These are two different processes in the OS. IP delivers the packet to the right process by looking at the IP header.

❖ What are the limits of the TCP's guaranteed service? It is said that TCP allows for reliable data transfer over an unreliable network layer, but what happens when there's a large network outage? i.e. is TCP explicit about what conditions it requires to guarantee reliable transfer of data?

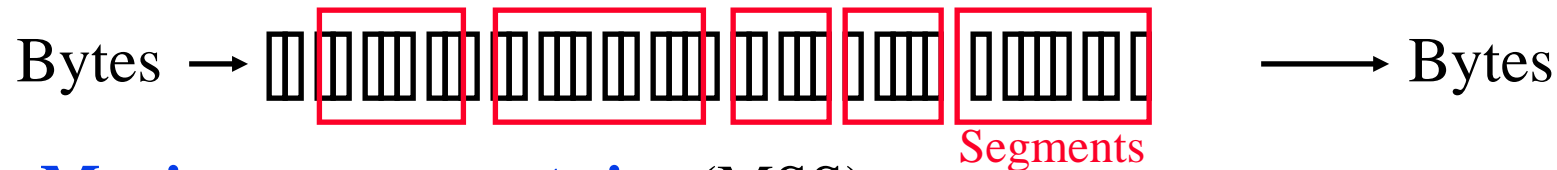
Yes, if there is a network outage, TCP will not deliver and so whatever is delivered is reliable.

Key Features of TCP

❑ **Point-to-Point:** One sender, one receiver

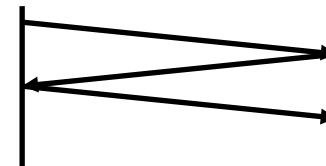


❑ **Byte Stream:** No message boundaries.
TCP makes “segments”

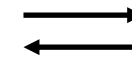


❑ **Maximum segment size (MSS)**

❑ **Connection Oriented:** Handshake to initialize states before data exchange

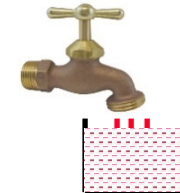


❑ **Full Duplex:** Bidirectional data flow in one connection



❑ **Reliable:** In-order byte delivery

❑ **Flow control:** To avoid receiver buffer overflow



❑ **Congestion control:** To avoid network router buffer overflow

Student Questions

❖ For review question R5 from the textbook, why does it say "voice and video traffic is often sent over TCP rather than UDP"? I thought that videos are more often sent with UDP.

The only reason, video traffic uses TCP is for the video parameters and control. The data itself is sent using UDP.

TCP

- ❑ Transmission Control Protocol
- ❑ Key Services:
 - **Send**: Please send when convenient
 - **Data stream push**: Destination TCP, please deliver it immediately to the receiving application.
⇒ Source TCP, please send it now.
Set on last packet of an application message.
 - **Urgent data signaling**: Destination TCP, please give this urgent data to the user out-of-band.
Generally used for CTRL-C.

Student Questions

❑ If push is used wouldn't that scramble the data in the stream?

No, everything is pushed.

❑ Those key features seem to require the cooperation of intermediate hop, but since we didn't implement the transport layer in routers, how does the router know is this urgent data or not?

Routers do not distinguish urgent data or non-urgent data.

❑ Can we also simulate TCP transmission on our localhost?

Yes. Two processes can send and receive via TCP. For example, a webserver and browser on the same host.

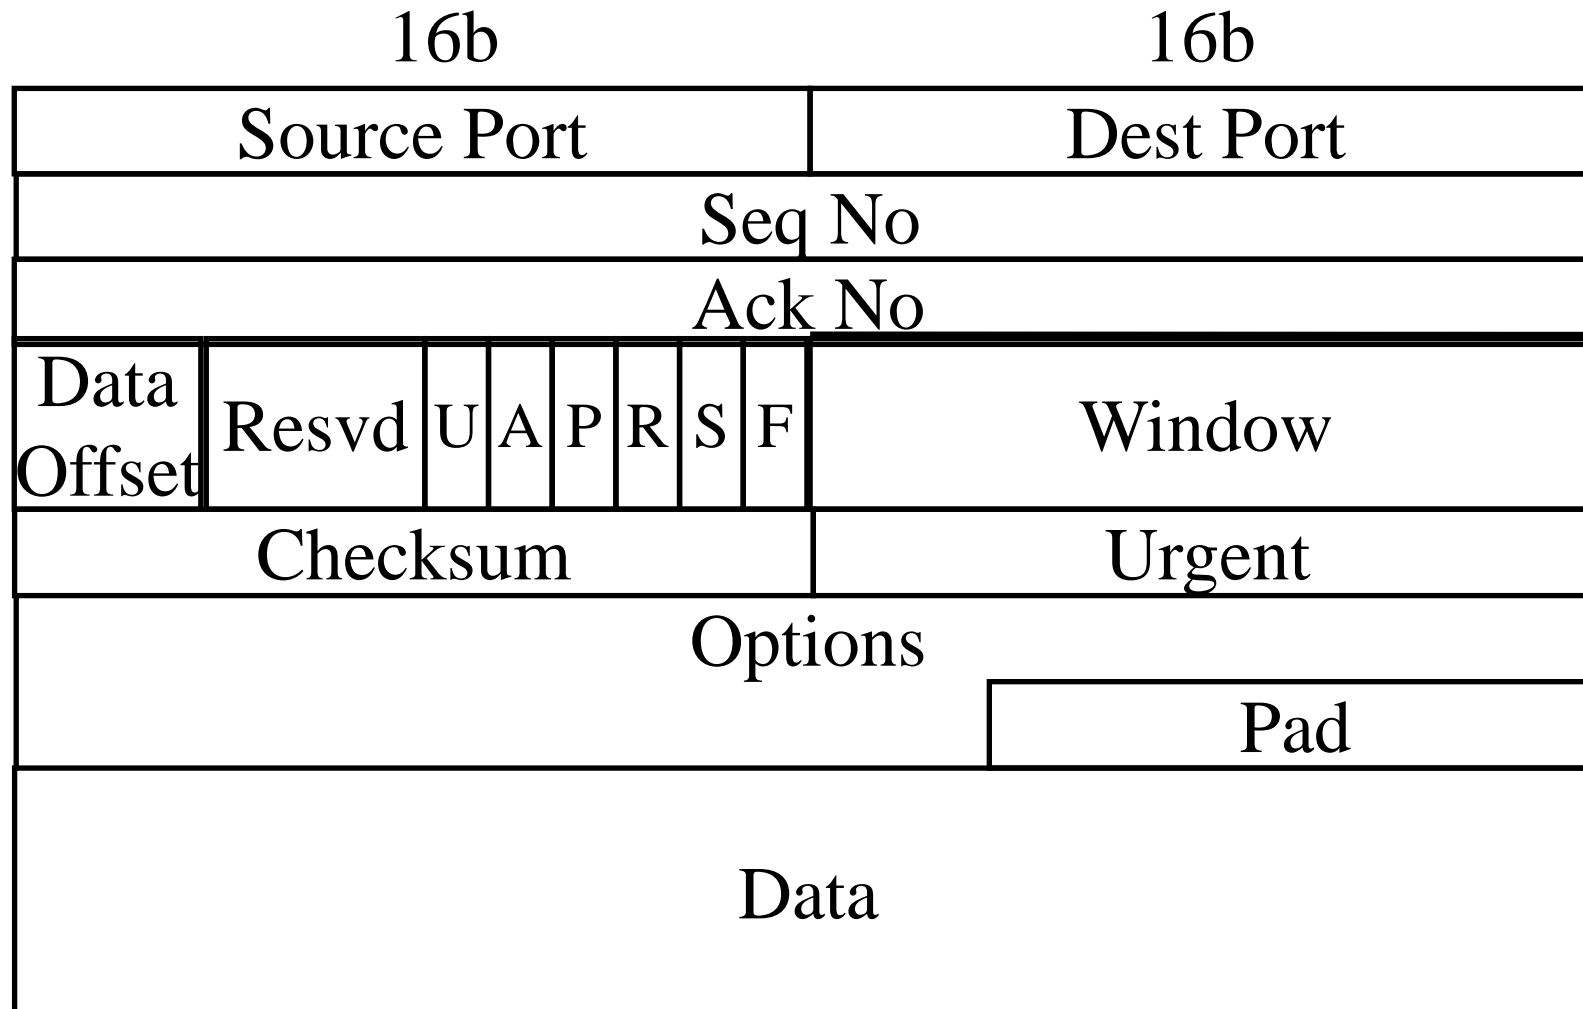
❑ When using send, how is convenience determined?

Convenience = Do I have enough or enough time has passed?

❑ If CTRL-C is pressed, does it mean the remaining unsent data will never be sent?

CTRL-C does not have any additional data associated with it. Other urgent tasks may have

TCP Segment Format (Cont)



Student Questions

What happens if I send more than 2^{32} segments of data? Would the sequence and acknowledgment numbers overflow?

The sequence number and acks "wrap around." This is not a problem if it happens after a large time.

Can you explain this again? It's hard to understand what you're pointing at because the video is a screen recording.

Sure.

Will there ever be a time that the reserved bits get assigned to be used for something, or are they just never going to be used?

Often. They use it for new features.

Is the Pad used to fill up 32 bits or a multiple of 32 bits?

The minimum number of bits required to make it a multiple of 32 bits.

You mentioned that there are 15 words at most. Where does this number come from?

Data offset is 4 bits. So max header = $2^4 - 1$

What is the header part in this format, everything before data? *Yes.*

TCP Header Fields

- ❑ **Source Port** (16 bits): Identifies source user process
- ❑ **Destination Port** (16 bits)
21 = FTP, 23 = Telnet, 53 = DNS, 80 = HTTP, ...
- ❑ **Sequence Number** (32 bits): Sequence number of the first byte in the segment. If SYN is present, this is the initial sequence number (ISN) and the first data byte is ISN+1.
- ❑ **Ack number** (32 bits): Next byte expected
- ❑ **Data offset** (4 bits): Number of 32-bit words in the header
- ❑ **Reserved** (6 bits)

Student Questions

❑ Could you explain why, if syn is present, the first data byte is ISN+1?

This is the first segment with data.

❑ Can we review what SYN refers to again?

SYN=Synchronize Sequence Number

❑ Are the reserved bits being held for future use?

Yes.

❑ Is the sender responsible for choosing the ISN?

Yes.

❑ Could you explain again what exactly the data offset is used for?

Offset=Where does the data begin?

❑ Why is it 16 bit?

If you are asking about port numbers, they could not imagine any more than 64K applications.

❑ Could you explain how ISN works and why the sequence number will be set to ISN+1?

ISN=Initial Sequence Number.

Every byte is numbered. The first byte is numbered ISN+1.

TCP Header (Cont)

- ❑ **Control** (6 bits): Urgent pointer field significant,
Ack field significant,
Push function,
Reset the connection,
Synchronize the sequence numbers,
No more data from sender



- ❑ **Window** (16 bits):
Will accept [Ack] to [Ack]+[window]-1

Student Questions

- ❑ What does it mean to synchronize the sequence numbers?
So the receiver know what byte numbers sender is going to send.

TCP Header (Cont)

- ❑ **Checksum** (16 bits): covers the segment plus a pseudo header. Includes the following fields from IP header: source and dest adr, protocol, segment length. Protects from IP misdelivery.
- ❑ **Urgent pointer** (16 bits): Points to the byte following urgent data. Lets receiver know how much data it should deliver right away out-of-band.
- ❑ **Options** (variable):
Max segment size (does not include TCP header, default 536 bytes), Window scale factor, Selective Ack permitted, Timestamp, No-Op, End-of-options

Student Questions

❑ Should we assume max segment size is 536 bytes if it comes up in the future?

Yes, iff not specified. Iff = If and only if

❑ Why is 536 bytes default MSS?
Can it be more?

MSS can be any number. To allow 512-byte blocks, they selected 536 which includes the application header. 9K and 64K are commonly used since file sizes are now large.

❑ Could you clarify why the TCP header has so many fields? Are all of them needed?

Yes. More have been added.

TCP Header (Cont)

- ❑ **Checksum** (16 bits): covers the segment plus a pseudo header. Includes the following fields from IP header: source and dest adr, protocol, segment length. Protects from IP misdelivery.
- ❑ **Urgent pointer** (16 bits): Points to the byte following urgent data. Lets receiver know how much data it should deliver right away out-of-band.
- ❑ **Options** (variable):
Max segment size (does not include TCP header, default 536 bytes), Window scale factor, Selective Ack permitted, Timestamp, No-Op, End-of-options

Student Questions

- ❖ This TCP checksum seems consider more fields than the checksum that takes only 16bits words before. Is the TCP checksum another checksum than before or it is just the full description now?

The TCP checksum is the same size as the checksum discussed before. It is the same checksum. There is no difference.

TCP Options

Kind	Length	Meaning
0	1	End of Valid options in header
1	1	No-op
2	4	Maximum Segment Size
3	3	Window Scale Factor
8	10	Timestamp

- ❑ **End of Options:** Stop looking for further option
- ❑ **No-op:** Ignore this byte. Used to align the next option on a 4-byte word boundary
- ❑ **Max Segment Size (MSS):** Does not include TCP header

Student Questions

❑ What is the pseudo-header?

A virtual header that includes some IP information. We will see it in Layer 3.

❑ Could you reexplain the chart? From my understanding, the kind is the type in the TLV. So if the type value is 3, then the length will be 3 and the meaning of the value is the window scale factor.

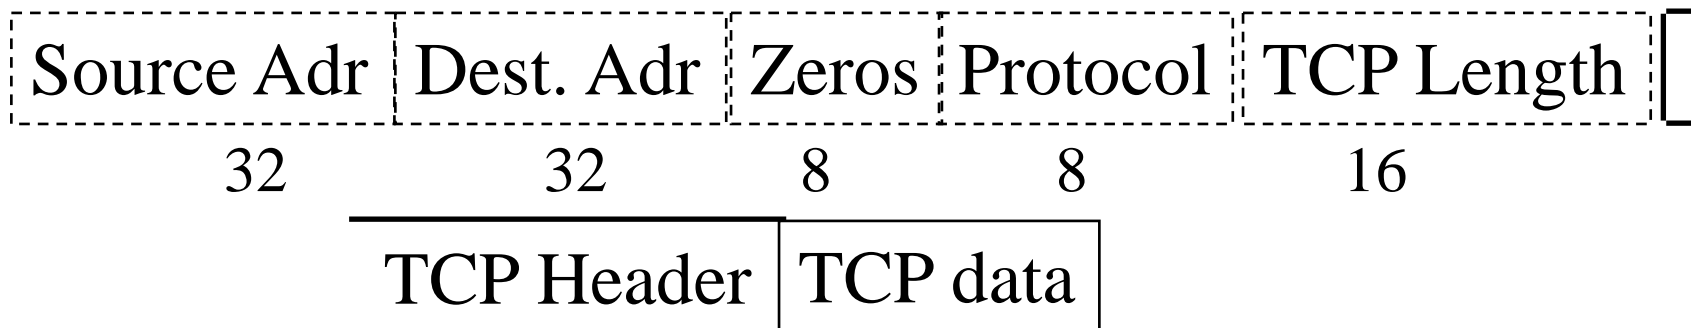
Correct.

❑ Can we interpret the table as the first column would be the type of options, each with a different number; the second column would be the length of the option, which will be in bytes.

Yes

TCP Checksum

- ❑ Checksum is the 16-bit one's complement of the one's complement sum of a pseudo header of information from the IP header, the TCP header, and the data, padded with zero octets at the end (if necessary) to make a multiple of two octets.
- ❑ Checksum field is filled with zeros initially
- ❑ TCP length (in octet) is not transmitted but used in calculations.
- ❑ Efficient implementation in RFC1071.



Student Questions

❑ Can you explain why the checksum is initially filled with zeros again?

Suppose A...B are message words. C is the checksum.

$$A+...+B+C=0$$

$$C=-(A+...+B)$$

At the source:

$$A+...+B+0 = -C$$

At the destination: _____ $A+...+B+C=0$

❑ Is there any difference between UDP Checksum and TCP Checksum?

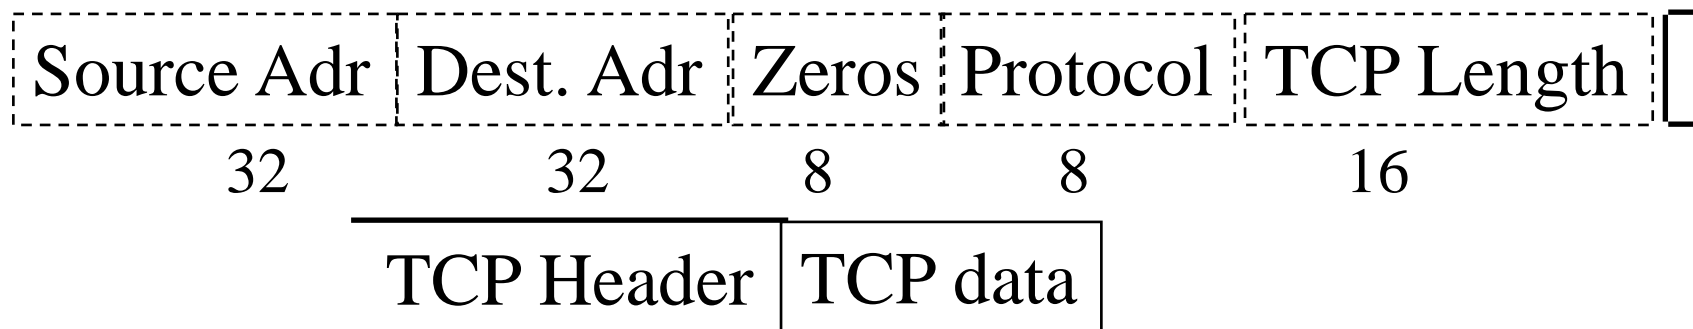
No

❑ In computing checksum, how do you add the source address to the protocol as they are of different bit lengths?

All IPv4 addresses are 32-bit long.

TCP Checksum

- ❑ Checksum is the 16-bit one's complement of the one's complement sum of a pseudo header of information from the IP header, the TCP header, and the data, padded with zero octets at the end (if necessary) to make a multiple of two octets.
- ❑ Checksum field is filled with zeros initially
- ❑ TCP length (in octet) is not transmitted but used in calculations.
- ❑ Efficient implementation in RFC1071.



Student Questions

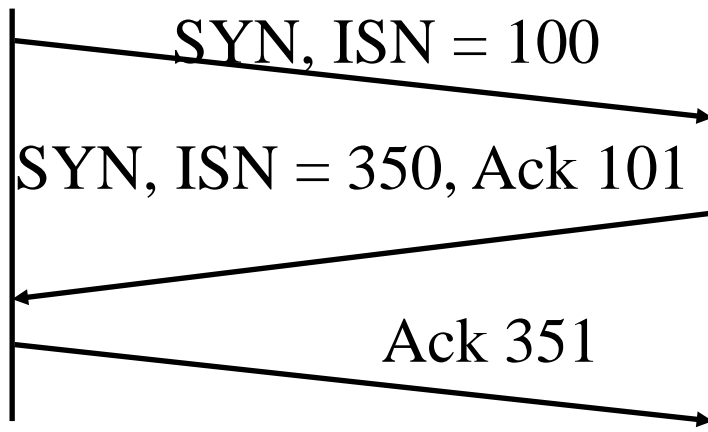
- ❖ The checksum for UDP was optional, but it's required for TCP, right?

Yes.

TCP Connection Management

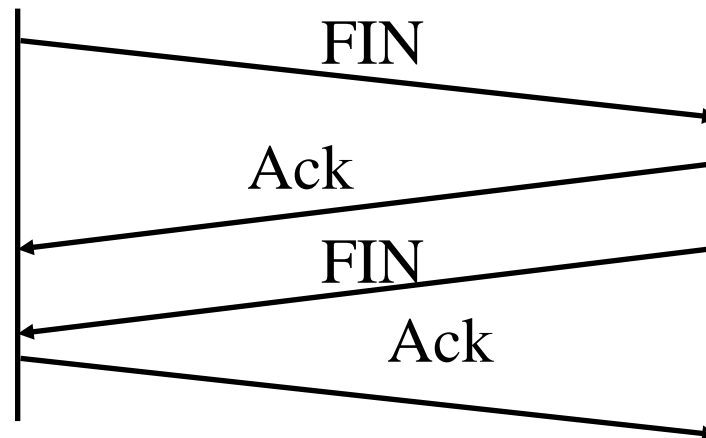
❑ Connection Establishment

- Three way handshake
- SYN flag set
⇒ Request for connection



❑ Connection Termination

- Close with FIN flag set
- Abort



Student Questions

❑ What settings must be configured to prevent against a SYN Flood Attack?

Start closing half-open connections as the resources decrease.

❑ When a client sends a SYN segment, do they also have to specify the port number to which they want to connect to?

Yes.

❑ If the first byte of a side (call this side A) starts at 101, why does the other side (call this side B) response "ack 101," which means everything up to 101, was received? Does this mean side A will send the bytes from 101 only after receiving "ark 101" from B?

"Ack 101" for TCP designers meant that I am expecting 101. I have received everything up to 100.

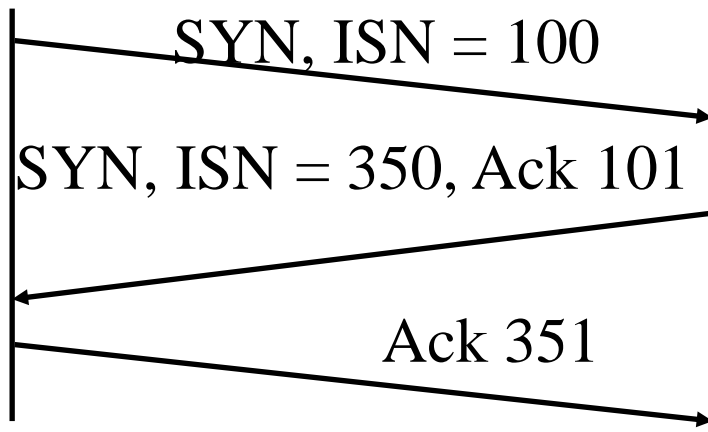
❑ Could you please go over the connection establishment and the importance of the SYN flag set?

SYN ⇒ The sequence number field contains the initial sequence number.

TCP Connection Management

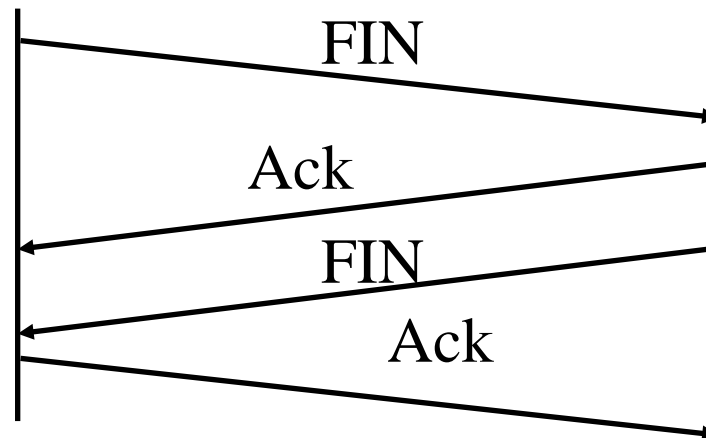
❑ Connection Establishment

- Three way handshake
- SYN flag set
⇒ Request for connection



❑ Connection Termination

- Close with FIN flag set
- Abort



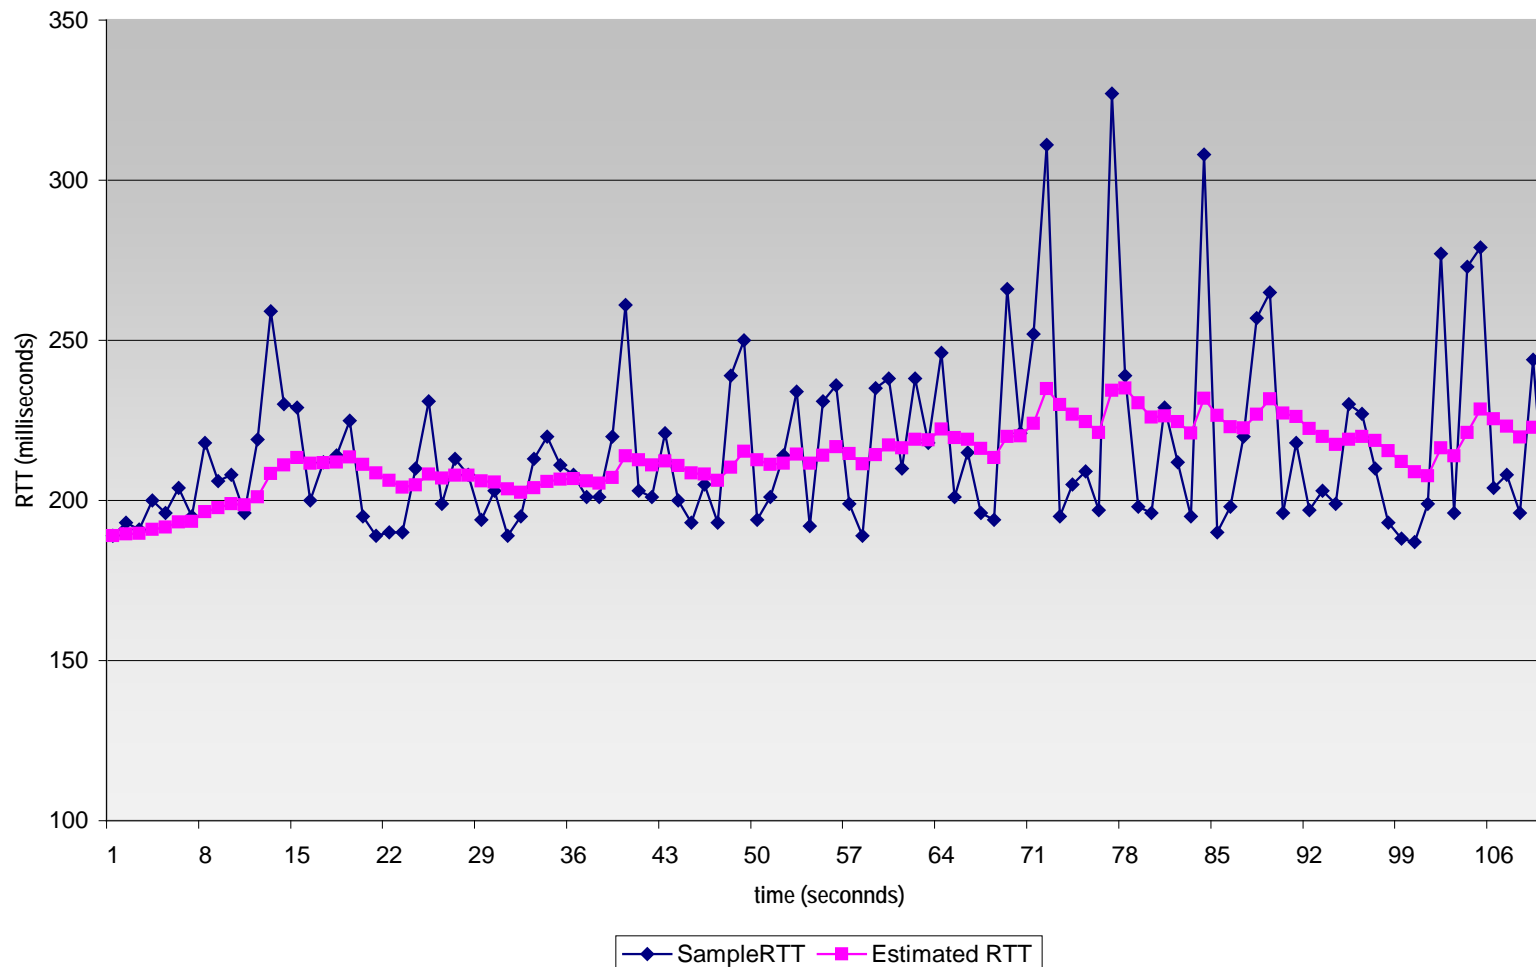
Student Questions

- ❖ What if the packet carrying FIN is lost, how will client detect this situation?

The source of FIN will timeout and retransmit.

Example RTT estimation:

RTT: gaia.cs.umass.edu to fantasia.eurecom.fr



Student Questions

❑ What is the relationship between RTT and propagation and transmission delay time?

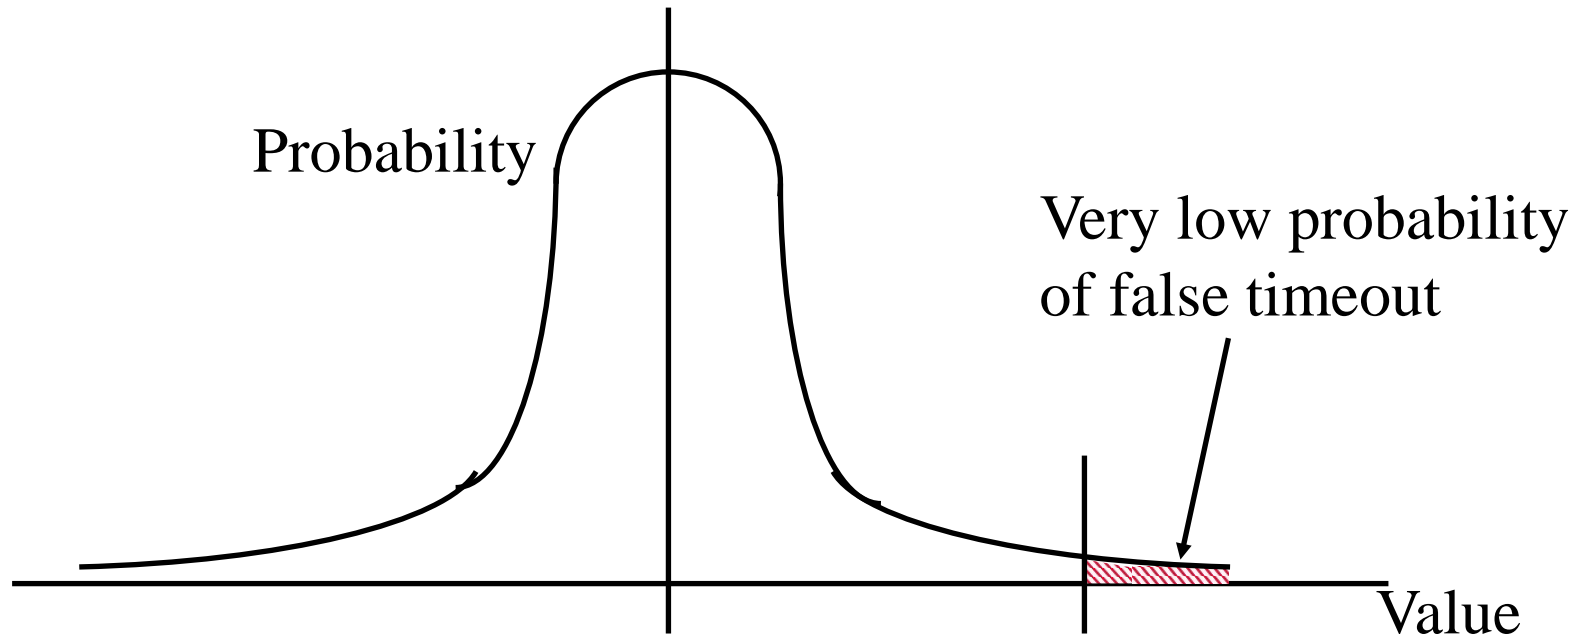
$RTT = 2 \times Propagation + Transmission + Queueing$

❑ In the video, you mention that you "don't have to send the data." Was this referring to the first or second Ack? Also, why is this the case?

Ack packets may or may not data going in the reverse direction.

Round Trip Time Estimation

- ❑ Measured round trip time (SampleRTT) is very random.
- ❑ $\text{EstimatedRTT} = (1 - \alpha)\text{EstimatedRTT} + \alpha \text{SampleRTT}$
- ❑ $\text{DevRTT} = (1 - \beta)\text{DevRTT} + \beta |\text{SampleRTT} - \text{EstimatedRTT}|$
- ❑ $\text{TimeoutInterval} = \text{EstimatedRTT} + 4 \text{DevRTT}$



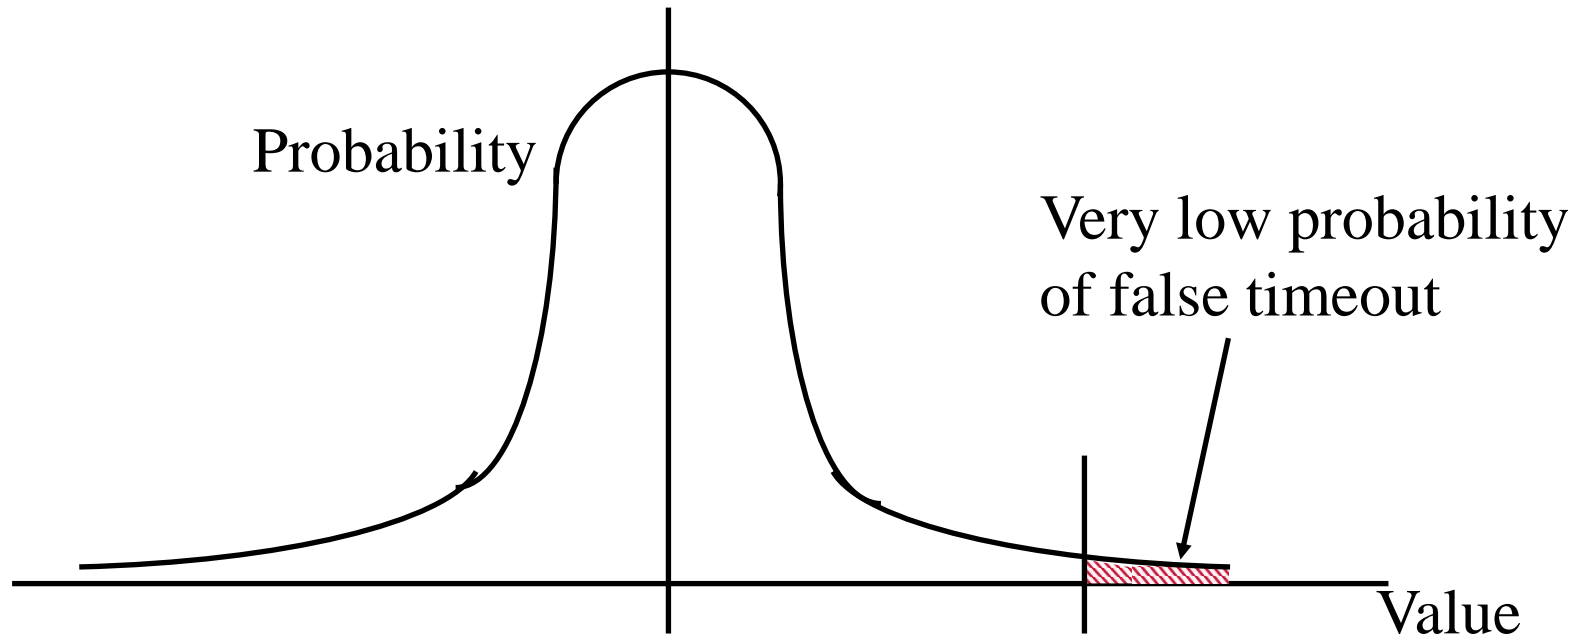
Student Questions

- ❑ The book says the recommended value of alpha is 1/8. Where does this come from/ who decides this?
Determined by experiments at the time.
- ❑ To clarify, this chart shows the probability of incorrectly measuring RTT/timeout based on the value of a and b?
Actual RTT is a random number. The red area shows the probability of RTT being more than Timeout ⇒ Packet was not lost but assumed lost

- ❑ Where do we get alpha and beta from?
Alpha and beta are parameters set by TCP users. Default values are normally used.
- ❑ Is RTT likely to follow a normal distribution?
Assumption
- ❑ Will there be any formula sheet on our exam that includes these?
Please see Video 0 on course

Round Trip Time Estimation

- ❑ Measured round trip time (SampleRTT) is very random.
- ❑ $\text{EstimatedRTT} = (1 - \alpha)\text{EstimatedRTT} + \alpha \text{SampleRTT}$
- ❑ $\text{DevRTT} = (1 - \beta)\text{DevRTT} + \beta |\text{SampleRTT} - \text{EstimatedRTT}|$
- ❑ $\text{TimeoutInterval} = \text{EstimatedRTT} + 4 \text{DevRTT}$



Student Questions

- ❖ Why sample RTT is random?
- ❖ In the round trip time estimation, what do " α " and " β " stand for?

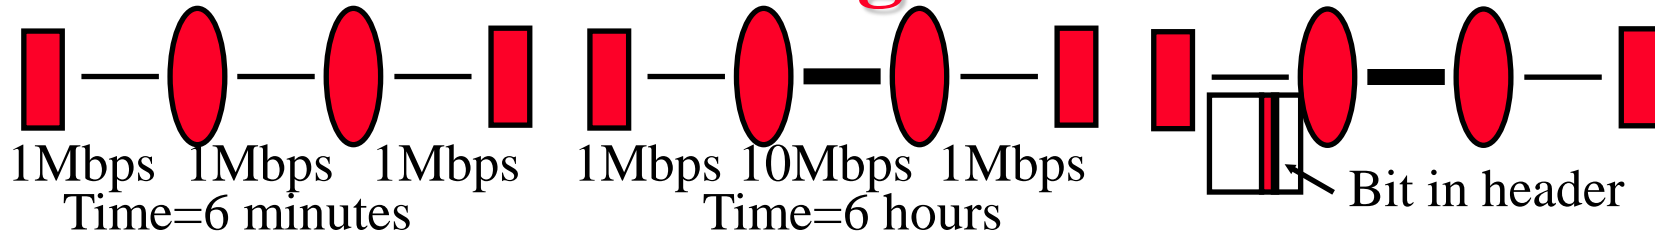
Because network traffic is random.

These are parameters selected by the network implementors or network managers. Must be between 0 and 1.

$\approx 1 \Rightarrow \text{Estimage} = \text{Sample}$. More random but current

$\approx 0 \Rightarrow \text{More stable but less current}$

Our Research on Congestion Control



- ❑ Early 1980s Digital Equipment Corporation (DEC) introduced Ethernet products
- ❑ Noticed that throughput goes down with a higher-speed link in middle (because no congestion mechanisms in TCP)
- ❑ Results:
 1. Timeout \Rightarrow Congestion
 \Rightarrow Reduce the TCP window to one on a timeout [Jain 1986]
 2. Routers should set a bit when congested (DECbit). [Jain, Ramakrishnan, Chiu 1988]
 3. Introduced the term “Congestion Avoidance”
 4. Additive increase and multiplicative decrease (AIMD principle) [Chiu and Jain 1989]
- ❑ There were presented to IETF in 1986.
 \Rightarrow Slow-start based on Timeout and AIMD [Van Jacobson 1988]

Student Questions

❑ Is the total speed always dependent on the link with the lowest speed?

Yes.

❖ Could you talk about congestion again?

Sure.

Slow Start Congestion Control

- ❑ Window = Flow control avoids receiver overrun
- ❑ Need congestion control to avoid network overrun
- ❑ The sender maintains two windows:
 - Credits from the receiver
 - Congestion window from the network
 - Congestion window is always less than the receiver window
- ❑ Starts with a congestion window (CWND) of 1 max segment size (MSS)
 - ⇒ Do not disturb existing connections too much.
- ❑ Increase CWND by 1 MSS every time an ack is received
- ❑ Assume CWND is in bytes

Student Questions

❑ You said multiple times that the packet size increases additively in "slow start", but your description only sounds exponential. Could you elaborate more on which ways the system is exponential and which ways it is additive?

Window is exponential in Time. Increase is additive

per ack $T=0$: $W=1$

In 1st RTT: 1 Ack $\Rightarrow W=2$

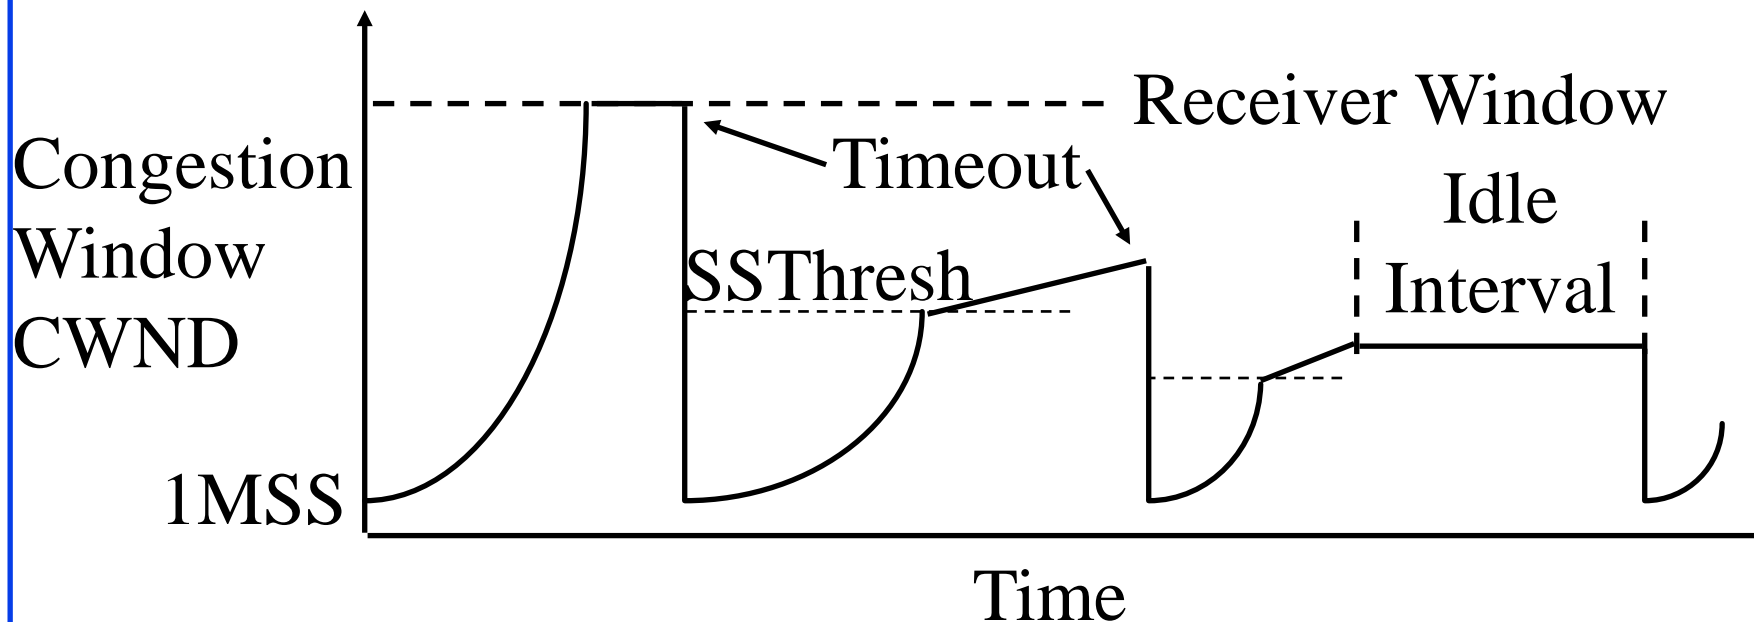
In 2nd RTT: 2 Acks $\Rightarrow W=4$ In 3rd RTT: 4 Acks $\Rightarrow W=8$

❑ Is window size controlled by the number of ACKs received?

No. It is controlled by the window size permitted by the receiver.

Slow Start (Cont)

- If segments lost, remember slow start threshold (SSThresh) to $CWND/2$
Set $CWND$ to 1 MSS
Increment by 1MSS per ack until SSThresh
Increment by $1 \text{ MSS} * \text{MSS} / CWND$ per ack afterwards



Student Questions

Slow Start (Cont)

- ❑ At the beginning, SSThresh = **Arbitrarily high value**
- ❑ After a long idle period (exceeding one round-trip time), reset the congestion window to one.
- ❑ If CWND is W MSS, W acks are received in one round trip.
- ❑ Below SSThresh, CWND is increased by 1MSS on every ack
 - ⇒ CWND increases to $2W$ MSS in one round trip
 - ⇒ CWND increases exponentially with timeExponential growth phase is also known as “**Slow start**” phase
- ❑ Above SSThresh, CWND is increased by $MSS/CWND$ on every ack
 - ⇒ CWND increases by 1 MSS in one round trip
 - ⇒ CWND increases linearly with timeThe linear growth phase is known as “**congestion avoidance**” phase

Student Questions

- ❑ What is meant by setting SSThresh initially to “arbitrarily high value”

I checked many different sources and here is the summary:

RFC 2001 - Dated Jan 1997 - says initially SSThresh should be set to 64kB.

This RFC is the first RFC describing Slow Start. However, it has been obsoleted by subsequent changes.

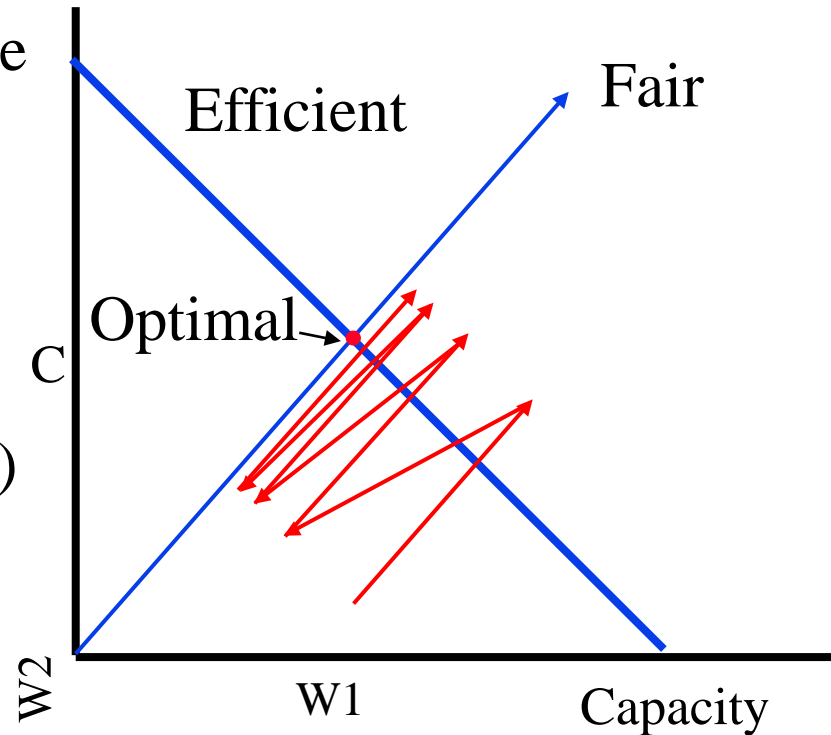
RFC 5681 - Dated Sept 2009 - says initially SSThresh should be set to arbitrarily high.

What changed in those 12 years? The networks became faster and 64kB, which was very high in 1997 became small and so no number was removed.

However, all documents state that the CWND can never exceed the Receiver Window. So slow start continues up to the receiver window if there is no loss. Setting to arbitrarily high value covers the case when the receiver window is changed over time. In this case, the slow start will continue until the point when CWND is equal to the receiver window.

AIMD Principle

- Additive Increase, Multiplicative Decrease
- $W1+W2 = \text{Capacity}$
⇒ Efficiency,
 $W1=W2 \Rightarrow \text{Fairness}$
- $(W1, W2)$ to $(W1+\Delta W, W2+\Delta W)$
⇒ Linear increase (45° line)
- $(W1, W2)$ to $(kW1, kW2)$
⇒ Multiplicative decrease
(line through origin)



Ref: D. Chiu and Raj Jain, "Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks," Journal of Computer Networks and ISDN, Vol. 17, No. 1, June 1989, pp. 1-14,
http://www.cse.wustl.edu/~jain/papers/cong_av.htm

Student Questions

- Could you explain the graph again as the video doesn't show where you're pointing on the board; will it always reach an equilibrium where efficiency and fairness intersect?

The optimal is the intersection of the fair and efficient line shown by the red dot.

Yes, AIMD will always take you to the optimal point.

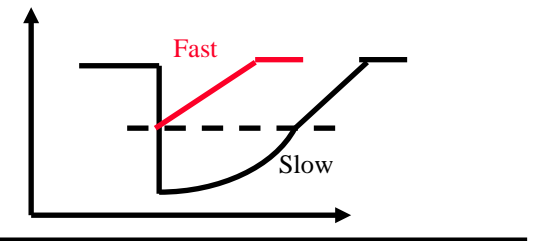
Fast Retransmit

- ❑ Optional – implemented in TCP Reno (Earlier version was TCP Tahoe)
- ❑ Duplicate Ack indicates a lost/out-of-order segment
- ❑ On receiving 3 duplicate acks (4th ack for the same segment):
 - Enter Fast Recovery mode
 - ❑ Retransmit missing segment
 - ❑ Set $SSThresh = CWND/2$
 - ❑ Set $CWND = SSThresh + 3 \text{ MSS}$ (**Note: CWND is inflated**)
 - ❑ Every subsequent duplicate ack: $CWND = CWND + 1 \text{ MSS}$
 - When a new ack (not a duplicate ack) is received
 - ❑ Exit fast recovery
 - ❑ Set $CWND = SSTHRESH$ (**Note: CWND is deflated back**)

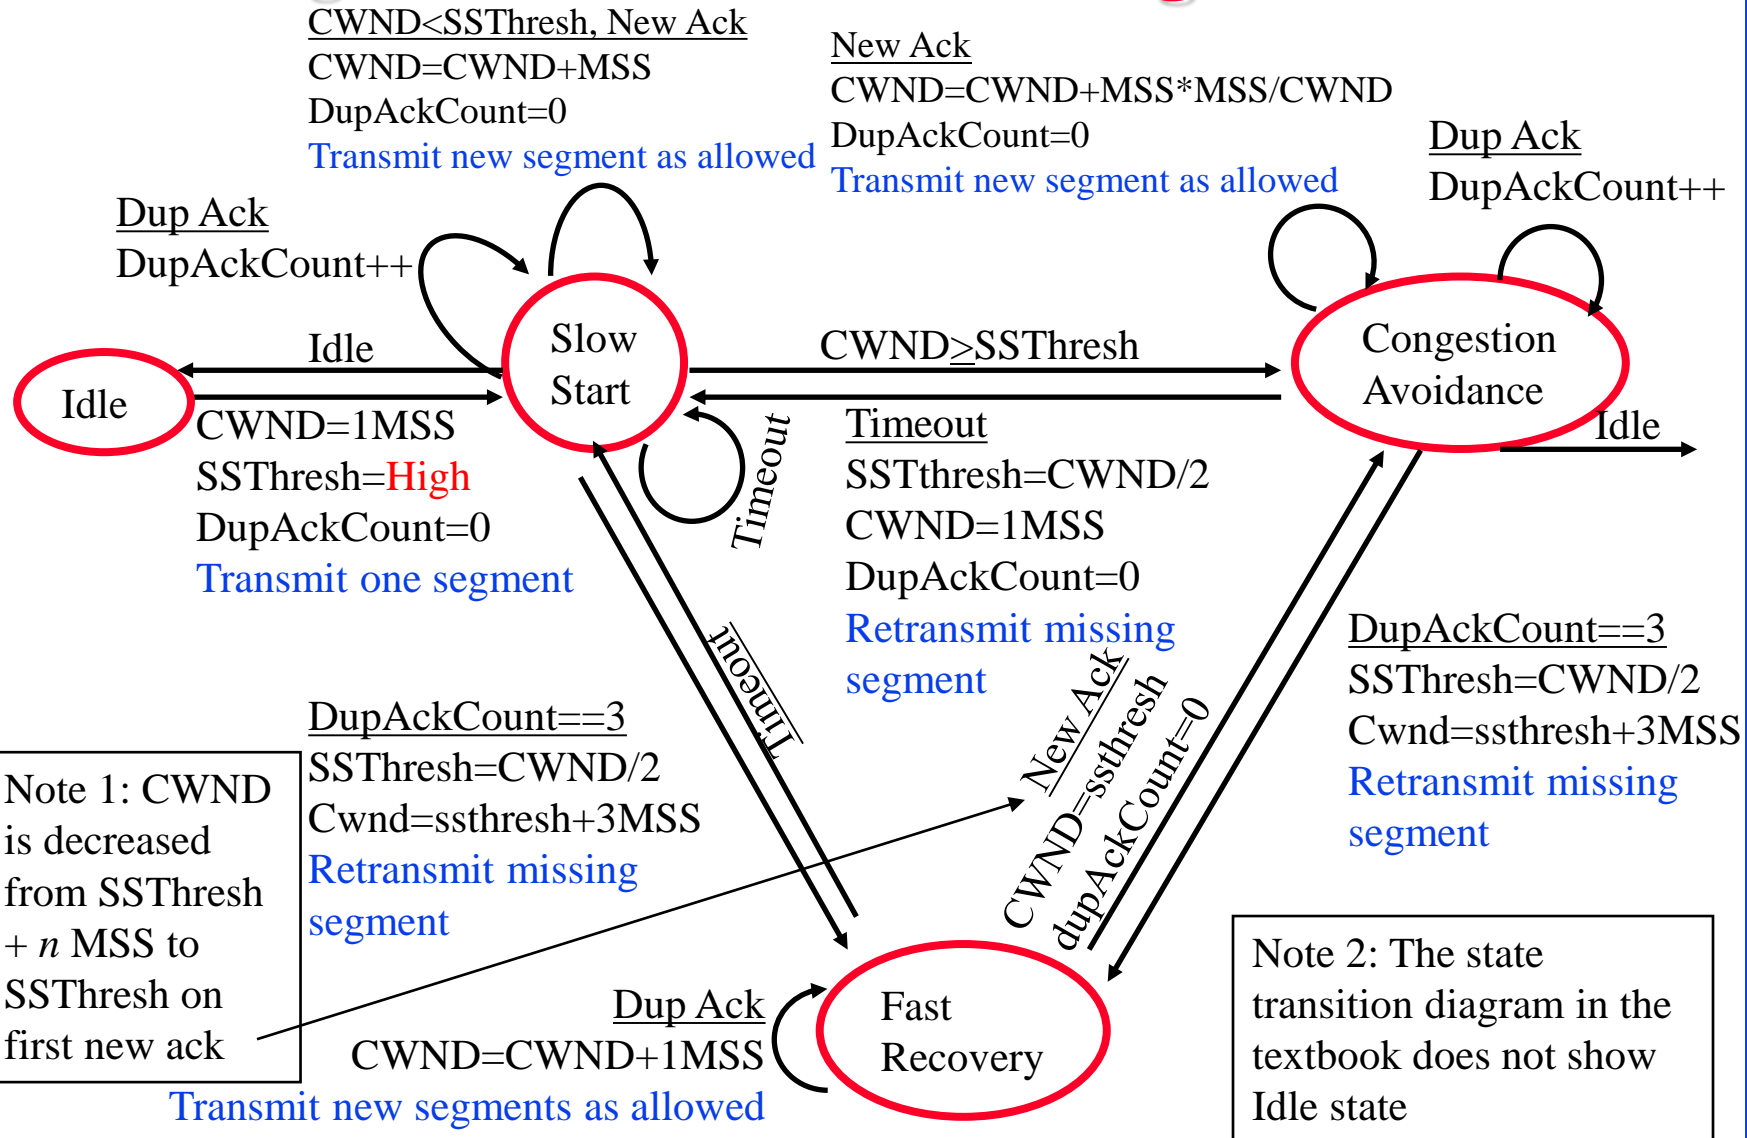
Student Questions

- ❑ What exactly is fast recovery doing? is it simply a transitional state of sending out the same data until an ack is finally received?

Fast recovery avoids the time loss due to dropping to one.



TCP Congestion Control State Diagram



Note 1: CWND is decreased from SSThresh + n MSS to SSThresh on first new ack

Note 2: The state transition diagram in the textbook does not show Idle state

Student Questions

- For the TCP Congestion Avoidance phase, for every new ACK, $cwnd = cwnd + MSS \cdot (MSS/cwnd)$. Graphs show the Congestion Avoidance phase to grow linearly, but is that a simplification? Because cwnd increases with every new ACK, the term $MSS \cdot (MSS/cwnd)$ shrinks with every ACK. Should the second term instead be $MSS \cdot (MSS/cwnd_0)$, where $cwnd_0$ is the initial value of cwnd when entering a round?

$T=0$ Window is W

$T=1RTT$ W acks are received.

Increase by $1/W$ per ack

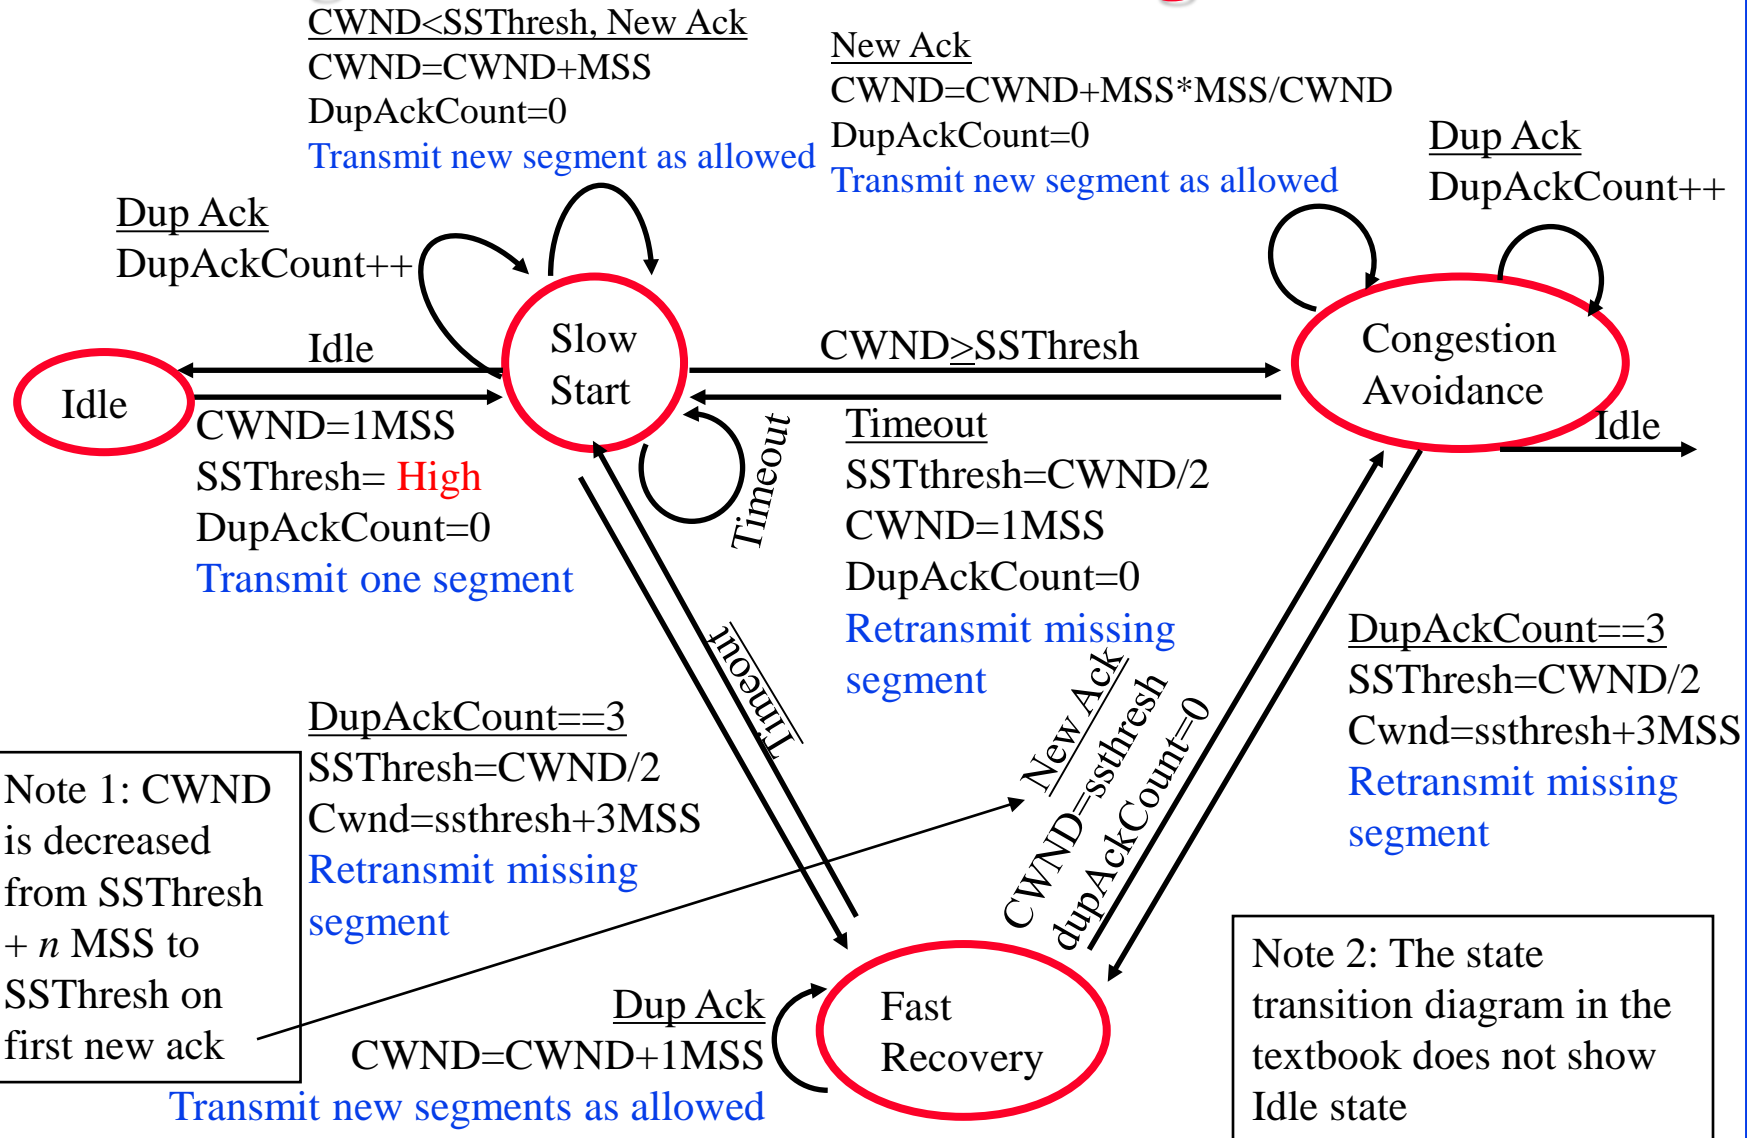
$W = W + W/W = W + 1$

So the increase is linear in time.

- Can you go back over this diagram one more time?

Sure.

TCP Congestion Control State Diagram



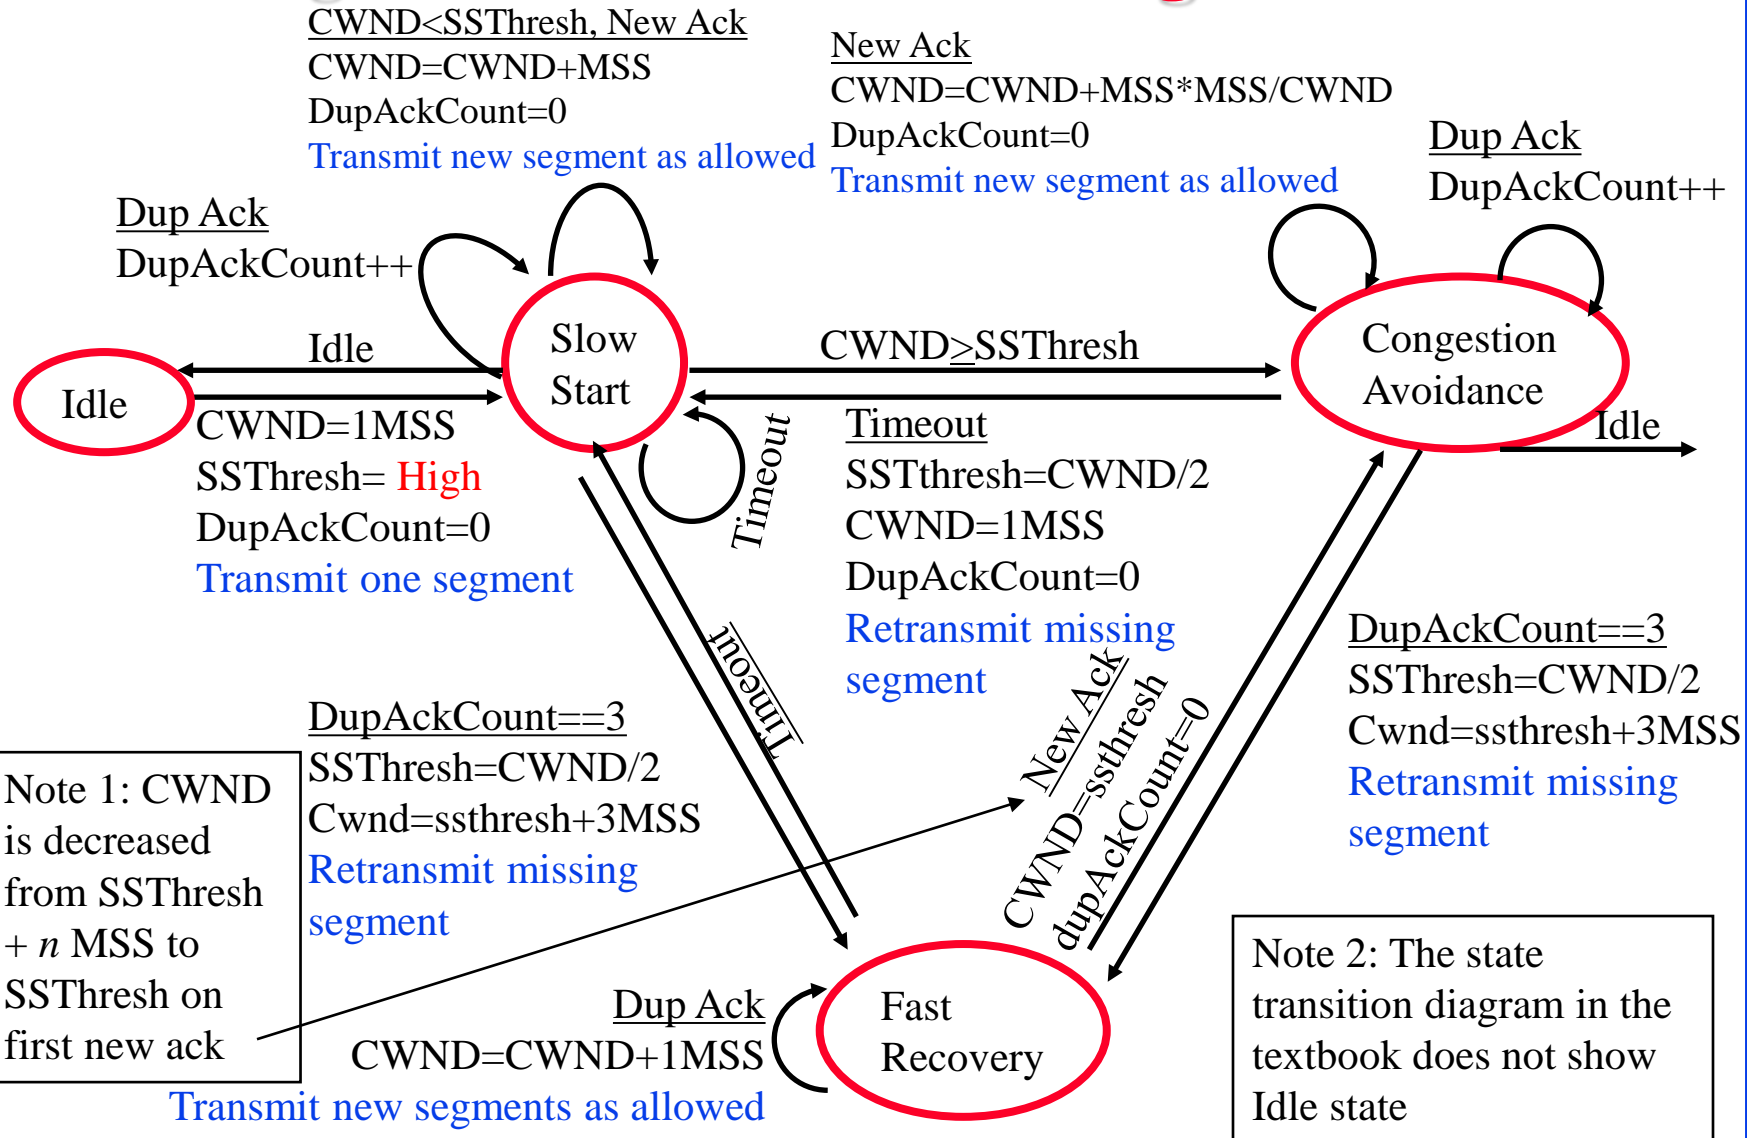
Note 1: CWND is decreased from SSThresh + n MSS to SSThresh on first new ack

Note 2: The state transition diagram in the textbook does not show Idle state

Student Questions

- ❖ Can we discuss the difference between a timeout and a triple duplicate ack again?
Timeouts are long. A number of “out-of-order” packets may reach the destination in that time. Each OoO packet is ack’ed with a “duplicate ack” and so triple duplicate ack indicates loss much sooner than timeout.
- ❖ If the last transmission round during slow start is CWND=16 (so that the next round is CWND=32) and if SSThresh is a number like 20, will it switch to congestion control with CWND=20 or switch with CWND=32.
CWND is increased after each ack. So it will switch as soon as CWND goes over 20.

TCP Congestion Control State Diagram



Note 1: CWND is decreased from SSThresh + n MSS to SSThresh on first new ack

Note 2: The state transition diagram in the textbook does not show Idle state

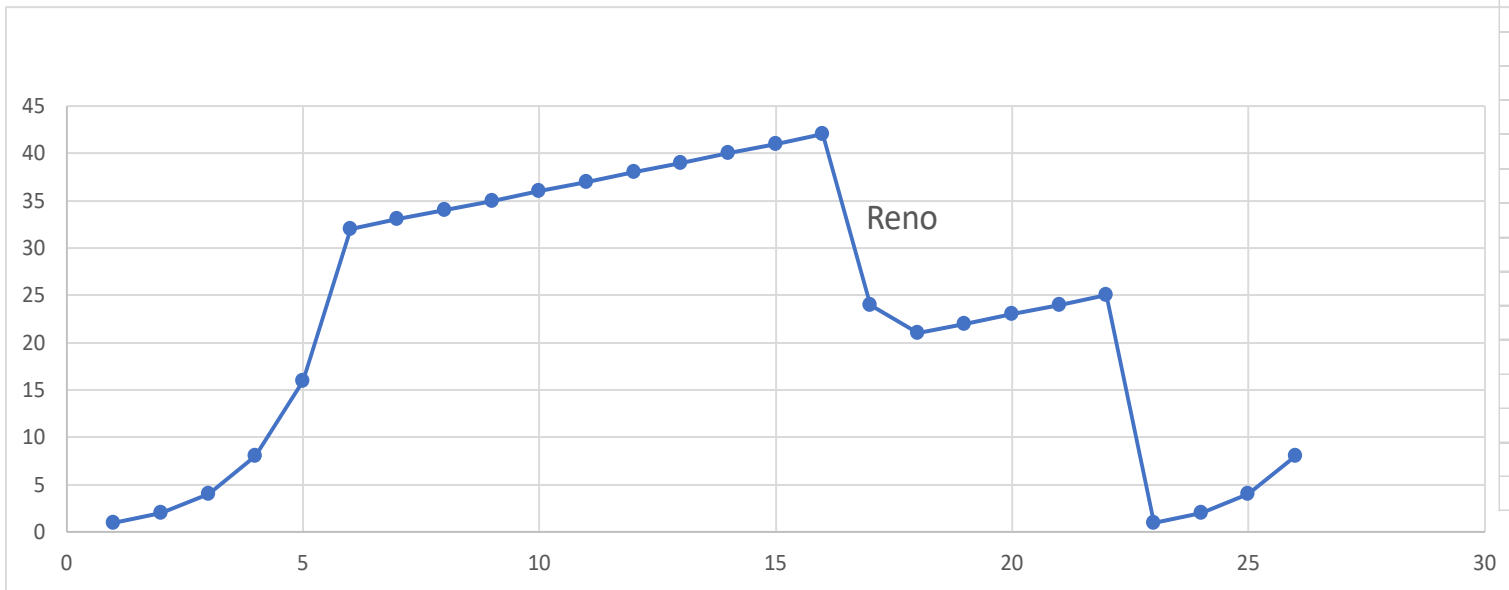
Student Questions

- ❖ if a package gets lost at the very beginning (obviously in slow start) due to random reasons, then it will switch to congestion avoidance after a while; but if the line is spare, then CWND will increase with a very slow rate, and it loses the advantage of slow start. Is it unavoidable?
Anything can be programmed provided you can show that it is better. Actually what is happening is the right thing to do. There is no need to avoid it.

Homework 3C: Slow Start

- [22 points] Consider Figure below. Assuming TCP Reno is the protocol experiencing the behavior shown above, answer the following questions. In all cases, you should provide a short discussion justifying your answer.

Round	CWND	Reno
1	1	1
2	2	2
3	4	4
4	8	8
5	16	16
6	32	32
7	33	33
8	34	34
9	35	35
10	36	36
11	37	37
12	38	38
13	39	39
14	40	40
15	41	41
16	42	42
17	24	24
18	21	21
19	22	22
20	23	23
21	24	24
22	25	25
23	1	1
24	2	2
25	4	4
26	8	8



Student Questions

- Can we go over homework 3C?
 - Can we get the answers to the homework?
-
- ❖ Could we go over the CWND graph problem?

Sure.

Homework 3C (Cont)

- ❑ A. Identify the interval of time when TCP slow start is operating.
- ❑ B. Identify the intervals of time when TCP congestion avoidance is operating.
- ❑ C. After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?
- ❑ D. After the 22nd transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?
- ❑ E. What is the initial value of ssthresh at the first transmission round?
- ❑ F. What is the value of ssthresh at the 18th transmission round?
- ❑ G. What is the value of ssthresh at the 24th transmission round?

Student Questions

Homework 3C (Cont)

- ❑ H. During what transmission round is the 70th segment sent?
- ❑ I. Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congestion window size and of ssthresh?
- ❑ J. Suppose TCP Tahoe is used (instead of TCP Reno), and assume that triple duplicate ACKs are received at the 16th round. What are the ssthresh and the congestion window size at the 19th round? (*Hint: You need to calculate CWND in 17-22nd rounds first. It will be different than that shown for Reno.*)
- ❑ K. Again suppose TCP Tahoe is used, and there is a timeout event at the end of 22nd round. How many packets have been sent out from 17th round till 22nd round, inclusive?

Student Questions

- ❑ Are we able to select the interval (idle interval) after which the connection breaks?

Yes. Applications set it according to their needs.

- ❑ Why SStresh to CWND/2 rather than CWND?

You need to reduce the load (window) during congestion.

- ❑ For clarification, If CWND is 15, would SStresh be 7 or 8?

In the code CWND is maintained in Bytes and so it is possible for it to be 7.5 MSS. However, when in this situation, the sender should not send a small segment of size 0.5MSS. This leads to issues. So the effective window size will be 7.

Ref: V. Jacobson and M. Karels, "Congestion Avoidance and Control" Nov. 1988,

<https://ee.lbl.gov/papers/congavoid.pdf>

This is a very slightly revised version of their SIGCOMM'88 paper.

TCP Average Throughput

□ Average Throughput =
$$\frac{1.22 \text{ MSS}}{\text{RTT} \sqrt{P}}$$

□ Here, P = Probability of Packet loss

□ Note 1: The formula is an approximation which does not apply at P=0 or P=1. At P=1, the throughput is zero. At P=0, the throughput is $\min\{1, (\text{Receiver Window}/\text{RTT})\}$

□ Note 1: The textbook uses L for probability of packet loss but it was used earlier for length of packets.

Student Questions

- So this equation only works when P <= 1%?
- Yes, it is an approximation which works at low non-zero values of P.
- As for the optional homework 3D, did we suppose to apply this average throughput formula for solving that question?

No. That homework is similar to slow start. You need to determine the window at each roundtrip and count the number of packets.

- Where does "1.22" come from?

This is curve-fitting.

- Is there a reason why we assume that packet loss cannot be greater than 10%?

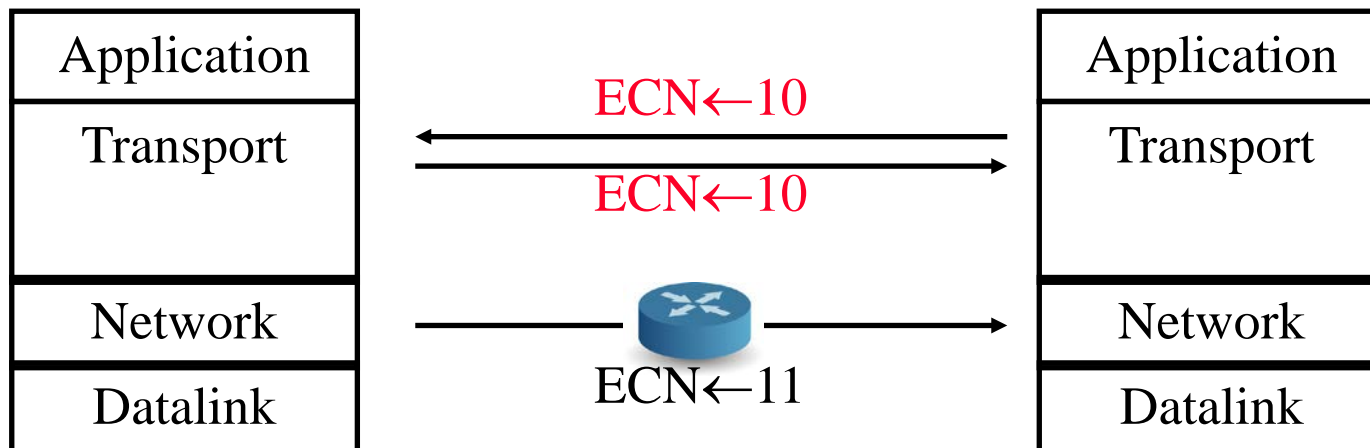
At 10%, every 10th packet will be lost and the network will seem very slow or broken.

- When the window lengths are different whether the different TCP connections are fair or not.

The user throughput depends upon the window size and RTT. Most congestion schemes are tested to check for fairness at the bottleneck even if RTTs are different. => Large windows for Large RTTs

Explicit Congestion Notification (ECN)

- Explicit congestion notification (ECN) is based on our DECbit research. Two bits in IP Header:
 - 00: Transport is not capable of ECN (e.g., UDP)
 - If non-zero, there are three possibilities:
 - 01: ECE = ECN Echo \Rightarrow Congestion was detected in the other direction
 - 10: CWR = Congestion window reduced
 - 11: CE = Congestion experienced
- When a router encounters congestion, instead of dropping the datagram, it marks the two bits as “11” congestion experienced



Student Questions

- Is just the IP / TCP is header getting sent back? Or is the whole packet itself getting sent back?

The bits are set in the TCP segments going in the reverse direction on the same TCP connection.

- How many packets in a row does it take before the router starts marking new ones as '11'?

If average queue length > 1. The routers are required to compute the average queue length. They can do this by exponentially weighted average.

- Are the ECN bits set in the Transport layer header or the Network layer header?

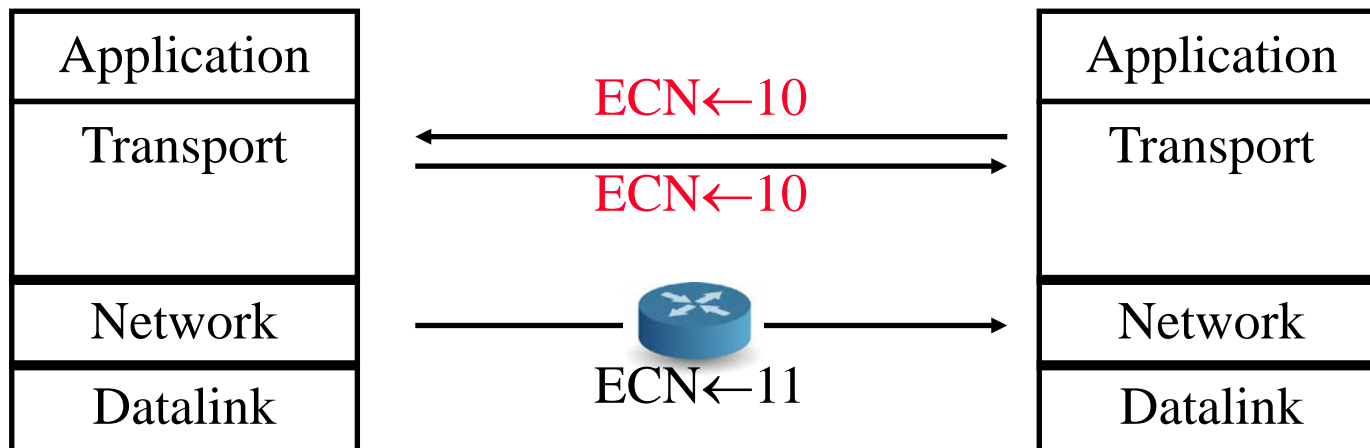
All of these bits are in network layer header.

- Can you please explain again how RTT and ECN are related? With a longer RTT and a congested network, would you not be sacrificing performance?

No. We are not sacrificing performance. We are maximizing performance. All of this was analyzed, measured, simulated before publishing the scheme.

Explicit Congestion Notification (ECN)

- Explicit congestion notification (ECN) is based on our DECbit research. Two bits in IP Header:
 - 00: Transport is not capable of ECN (e.g., UDP)
 - **If non-zero, there are three possibilities:**
 - **01: ECE = ECN Echo** ⇒ Congestion was detected in the other direction
 - **10: CWR = Congestion window reduced**
 - **11: CE = Congestion experienced**
 - When a router encounters congestion, instead of dropping the datagram, it marks the two bits as “11” congestion experienced



Student Questions

- ❖ Is congestion declared when the buffer of the router is overflowed or about to overflow?

In ECN congestion is declared when the average queue length is more than 1.

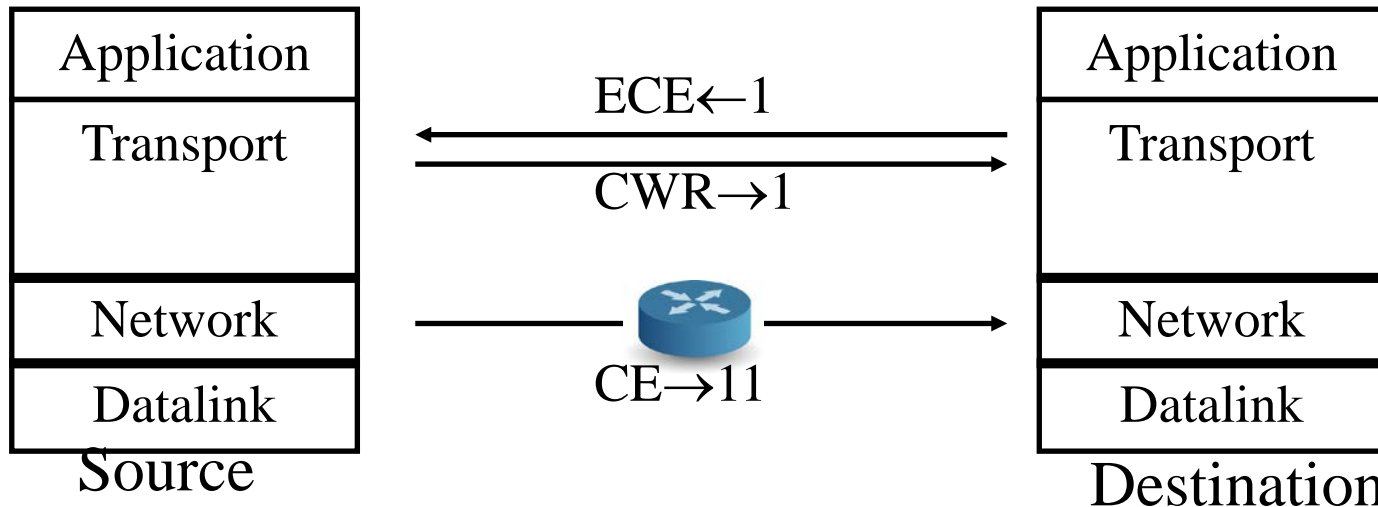
- Could you briefly review ECN (Slide 3-51)?

Sure.

- What's the difference between ECN and CE? I noticed them both send from source to destination.
- *ECN is the name of the field. CE is one possible value for that field.*

ECN (Cont)

- ❑ ECN uses two bits in TCP header: ECE and CWR
- ❑ On receiving “CE” code point, the receiver sends “ECN Echo (ECE)” flag in the TCP header
- ❑ On seeing the ECE flag, the source reduces its congestion window, and sets “Congestion Window Reduced (CWR) flag in outgoing segment
- ❑ On receiving “CWR” flag, the receiver, stops setting ECE bit



Ref: https://en.wikipedia.org/wiki/Explicit_Congestion_Notification

Student Questions

- ❑ What does CRW mean?

CRW= “I heard you. I have reduced the window”

- ❑ Is this just another protocol to manage congestion? Do all TCP protocols have different ways to manage congestion?

All routers implement all schemes. Some schemes are not implemented by some routers and so the schemes have to allow for such routers.

- ❑ ECN marks the two bits such as 00,01,10,11. Is the above situation happen at the same time?

No. At one instance only one marking is used.

- ❑ In the diagram, why are the "ECE" and "CWR" being transmitted? I thought it only gets transmitted during "CE."

ECE and CWR are always there but may be 0.

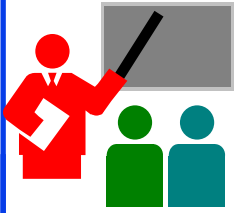
- ❑ Since ECN involves routers, is it in the network layer?

Yes.

- ❑ How is congestion control different from flow control?

Flow control is the control by the receiver.

Congestion control is by the network.



TCP: Summary

1. TCP uses **port numbers** for multiplexing
2. TCP provides reliable **full-duplex** connections.
3. TCP is **stream** based and has **window flow control**
4. **Slow-start congestion control** works on timeout
5. **Explicit congestion notification** works using ECN bits

Please see Slide 3-69 for discussion on CUBIC.

Read Sections 3.5, 3.6, and 3.7. Do R14-R19, P25, P26, P31, P32, P33, P40

Student Questions

- ❑ From Book Page 252: TCP has similarities with GBN/selective repeat, but does not explicitly use these, correct?

Selective ack has now been added to TCP. Particularly, helpful for highspeed networks. In SA, you can list the missing segments. So better than both GBN and SR.

- ❑ How does IP realizes the congestion?

By watching its queues.

- ❑ Why is ECN not possible if both bits are set, or neither are set?

ECN is always possible. If both bits are clear, there is no congestion, if both bits are 1, there is congestion.

- ❑ Could you briefly explain again what it means to reduce the congestion window?

Multiplicative decrease by a factor.

- ❑ What if the Header field for ECN is corrupted and sends the fake information to the source?

The source will reduce the window.

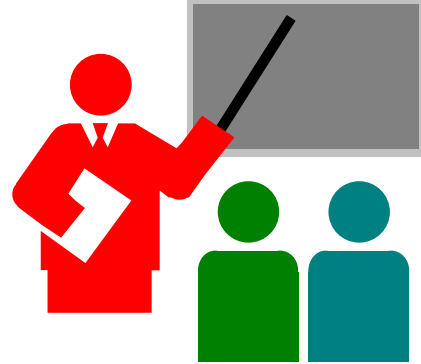
- ❑ Does TCP have any way to fight against UDP hogging bandwidth?

Indirectly. The queue length includes both kinds of packets.

- ❑ How would the network know if the congestion has been resolved? What flags would be set?

CE, ECE, and CRA will be zero.

Summary



1. **Multiplexing/demultiplexing** by a combination of source and destination IP addresses and port numbers.
2. **Longer distance or higher speed**
⇒ Larger α ⇒ Larger window is better
3. Window flow control is better for long-distance or high-speed networks
4. UDP is connectionless and simple.
No flow/error control. Has error **detection.**
5. TCP provides **full-duplex** connections with flow/error/**congestion** control.

Student Questions

- ❑ Shouldn't TCP be using both port numbers and IP addresses for multiplexing?

IP takes care of IP address. TCP takes care of port #.

- ❑ Can we say that TCP is a layer of abstraction on top of UDP?

You can implement that way but most people think of them as side-by-side and not one on the top of another.

Lab 3: Reliable Transport Protocol

[60 points] Overview

In this laboratory programming assignment, you will be writing the sending and receiving transport-level code for implementing a simple reliable data transfer protocol. There are two versions of this lab, the Alternating-Bit-Protocol version and the Go-Back-N version. This lab should be **fun** since your implementation will differ very little from what would be required in a real-world situation.

Since you probably don't have standalone machines (with an OS that you can modify), your code will have to execute in a simulated hardware/software environment. However, the programming interface provided to your routines, i.e., the code that would call your entities from above and from below is very close to what is done in an actual UNIX environment. (Indeed, the software interfaces described in this programming assignment are much more realistic than the infinite loop senders and receivers that many texts describe). Stopping/starting of timers are also simulated, and timer interrupts will cause your timer handling routine to be activated.

The routines you will write

The procedures you will write are for the sending entity (A) and the receiving entity (B). Only unidirectional transfer of data (from A to B) is required. Of course, the B side will have to send packets to A to acknowledge (positively or negatively) receipt of data. Your routines are to be implemented in the form of the procedures described below. These procedures will be called by (and will call) procedures that I have written which emulate a network environment. The overall structure of the environment is shown in [Figure Lab.3-1](#) (structure of the emulated environment):

The unit of data passed between the upper layers and your protocols is a *message*, which is declared as:

```
struct msg { char data[20];  
};
```

This declaration, and all other data structure and emulator routines, as well as stub routines (i.e., those you are to complete) are in the file, **prog2.c** (<http://gaia.cs.umass.edu/kurose/transport/prog2.c>). Your sending entity will thus receive data in 20-byte chunks from layer5; your receiving entity should deliver 20-byte chunks of correctly received data to layer5 at the receiving side.

Student Questions

- Is the deadline for Lab3 on 16th or 14th? *16th.*
- The alternating bits version is described, but the Go-Back-N version is not. Can you explain it a bit?

Go-back-N was described in Slide 3-21.

- Where are the instructions for the go-back N version?

Removed to reduce student load.

Could you elaborate more about the implementation of the GBN version?

- Is the design document required for this lab?

No.

- Is the lab also written in Python? *Most of the system software uses C. Python is used for applications. TCP is a part of the operating systems.*

Lab 3 (Cont)

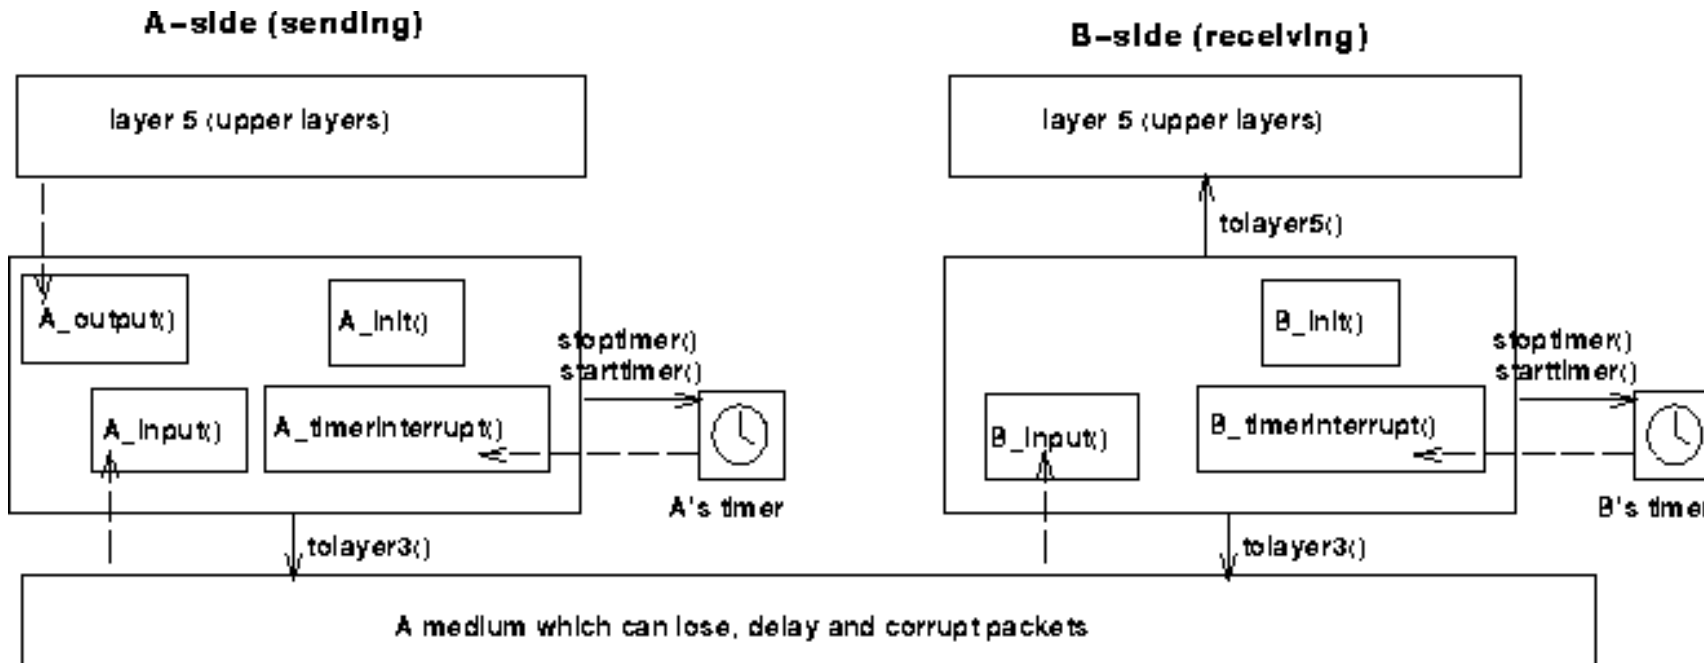


Figure Lab.3-1

Student Questions

- ❖ When we do lab 3 for GBN, if the packet received by the receiver from the application layer is outside the window, or meets buffer overflow issues, how to fix it?

There is no flow control on the traffic coming from the application to TCP, packets can be lost due to buffer overflow. Inter-process communication in the OS may prevent the application from proceeding since this is local.

Lab 3 (Cont)

The unit of data passed between your routines and the network layer is the *packet*, which is declared as:

```
struct pkt { int seqnum; int acknum;  
int checksum; char payload[20];  
};
```

Your routines will fill in the payload field from the message data passed down from layer5. The other packet fields will be used by your protocols to insure reliable delivery, as we've seen in class.

The routines you will write are detailed below. As noted above, such procedures in real-life would be part of the operating system, and would be called by other procedures in the operating system.

A_output(message), where message is a structure of type msg, containing data to be sent to the B-side. This routine will be called whenever the upper layer at the sending side (A) has a message to send. It is the job of your protocol to insure that the data in such a message is delivered in-order, and correctly, to the receiving side upper layer.

A_input(packet), where packet is a structure of type pkt. This routine will be called whenever a packet sent from the B-side (i.e., as a result of a tolayer3() being done by a B-side procedure) arrives at the A-side. packet is the (possibly corrupted) packet sent from the B-side.

A_timerinterrupt() This routine will be called when A's timer expires (thus generating a timer interrupt). You'll probably want to use this routine to control the retransmission of packets. See starttimer() and stoptimer() below for how the timer is started and stopped.

A_init() This routine will be called once, before any of your other A-side routines are called. It can be used to do any required initialization.

B_input(packet), where packet is a structure of type pkt. This routine will be called whenever a packet sent from the A-side (i.e., as a result of a tolayer3() being done by a A-side procedure) arrives at the B-side. packet is the (possibly corrupted) packet sent from the A-side.

B_init() This routine will be called once, before any of your other B-side routines are called. It can be used to do any required initialization.

Student Questions

Lab 3 (Cont)

Software Interfaces

The procedures described above are the ones that you will write. I have written the following routines which can be called by your routines:

starttimer(calling_entity,increment), where *calling_entity* is either 0 (for starting the A-side timer) or 1 (for starting the B side timer), and *increment* is a *float* value indicating the amount of time that will pass before the timer interrupts. A's timer should only be started (or stopped) by A-side routines, and similarly for the B-side timer. To give you an idea of the appropriate increment value to use: a packet sent into the network takes an average of 5 time units to arrive at the other side when there are no other messages in the medium.

stoptimer(calling_entity), where *calling_entity* is either 0 (for stopping the A-side timer) or 1 (for stopping the B side timer).

tolayer3(calling_entity,packet), where *calling_entity* is either 0 (for the A-side send) or 1 (for the B side send), and *packet* is a structure of type *pkt*. Calling this routine will cause the packet to be sent into the network, destined for the other entity.

tolayer5(calling_entity,message), where *calling_entity* is either 0 (for A-side delivery to layer 5) or 1 (for B-side delivery to layer 5), and *message* is a structure of type *msg*. With unidirectional data transfer, you would only be calling this with *calling_entity* equal to 1 (delivery to the B-side). Calling this routine will cause data to be passed up to layer 5.

Student Questions

Lab 3 (Cont)

The simulated network environment

A call to procedure `tolayer3()` sends packets into the medium (i.e., into the network layer). Your procedures `A_input()` and `B_input()` are called when a packet is to be delivered from the medium to your protocol layer.

The medium is capable of corrupting and losing packets. It will not reorder packets. When you compile your procedures and my procedures together and run the resulting program, you will be asked to specify values regarding the simulated network environment:

Number of messages to simulate. My emulator (and your routines) will stop as soon as this number of messages have been passed down from layer 5, regardless of whether or not all of the messages have been correctly delivered. Thus, you need **not** worry about undelivered or unACK'ed messages still in your sender when the emulator stops. Note that if you set this value to 1, your program will terminate immediately, before the message is delivered to the other side. Thus, this value should always be greater than 1.

Loss. You are asked to specify a packet loss probability. A value of 0.1 would mean that one in ten packets (on average) are lost.

Corruption. You are asked to specify a packet loss probability. A value of 0.2 would mean that one in five packets (on average) are corrupted. Note that the contents of payload, sequence, ack, or checksum fields can be corrupted. Your checksum should thus include the data, sequence, and ack fields.

Tracing. Setting a tracing value of 1 or 2 will print out useful information about what is going on inside the emulation (e.g., what's happening to packets and timers). A tracing value of 0 will turn this off. A tracing value greater than 2 will display all sorts of odd messages that are for my own emulator-debugging purposes. A tracing value of 2 may be helpful to you in debugging your code. You should keep in mind that *real* implementors do not have underlying networks that provide such nice information about what is going to happen to their packets!

Average time between messages from sender's layer5. You can set this value to any non-zero, positive value. Note that the smaller the value you choose, the faster packets will be arriving to your sender.

Student Questions

Lab 3 (Cont)

The Alternating-Bit-Protocol Version of this lab.

You are to write the procedures, `A_output()`, `A_input()`, `A_timerinterrupt()`, `A_init()`, `B_input()`, and `B_init()` which together will implement a stop-and-wait (i.e., the alternating bit protocol, which we referred to as rdt3.0 in the text) unidirectional transfer of data from the A-side to the B-side. **Your protocol should use both ACK and NACK messages.**

You should choose a very large value for the average time between messages from sender's layer5, so that your sender is never called while it still has an outstanding, unacknowledged message it is trying to send to the receiver. I'd suggest you choose a value of 1000. You should also perform a check in your sender to make sure that when `A_output()` is called, there is no message currently in transit. If there is, you can simply ignore (drop) the data being passed to the `A_output()` routine.

You should put your procedures in a file called `prog2.c`. You will need the initial version of this file, containing the emulation routines we have written for you, and the stubs for your procedures. You can obtain this program from <http://gaia.cs.umass.edu/kurose/transport/prog2.c>.

This lab can be completed on any machine supporting C. It makes no use of UNIX features. (You can simply copy the `prog2.c` file to whatever machine and OS you choose).

We recommend that you should hand in a code listing, a design document, and sample output. For your sample output, your procedures might print out a message whenever an event occurs at your sender or receiver (a message/packet arrival, or a timer interrupt) as well as any action taken in response. You might want to hand in output for a run up to the point (approximately) when 10 messages have been ACK'ed correctly at the receiver, a loss probability of 0.1, and a corruption probability of 0.3, and a trace level of 2. You might want to annotate your printout with a colored pen showing how your protocol correctly recovered from packet loss and corruption.

Note 1: The code requires GCC 4.8.

Ubuntu 14.0.4 comes with GCC 4.8. So you may need to install Ubuntu 14.0.4 in a virtual machine.

Note 2: Some students have suggested to add the line “`#include <stdlib.h>`” and removing all instances of “`char *malloc()`.”

Student Questions

Lab 3 (Cont)

Helpful Hints and the like

Checksumming. You can use whatever approach for checksumming you want. Remember that the sequence number and ack field can also be corrupted. We would suggest a TCP-like checksum, which consists of the sum of the (integer) sequence and ack field values, added to a character-by-character sum of the payload field of the packet (i.e., treat each character as if it were an 8 bit integer and just add them together).

Note that any shared "state" among your routines needs to be in the form of global variables. Note also that any information that your procedures need to save from one invocation to the next must also be a global (or static) variable. For example, your routines will need to keep a copy of a packet for possible retransmission. It would probably be a good idea for such a data structure to be a global variable in your code. Note, however, that if one of your global variables is used by your sender side, that variable should **NOT** be accessed by the receiving side entity, since in real life, communicating entities connected only by a communication channel can not share global variables.

There is a float global variable called *time* that you can access from within your code to help you out with your diagnostics msgs.

START SIMPLE. Set the probabilities of loss and corruption to zero and test out your routines. Better yet, design and implement your procedures for the case of no loss and no corruption, and get them working first. Then handle the case of one of these probabilities being non-zero, and then finally both being non-zero.

Debugging. We'd recommend that you set the tracing level to 2 and put LOTS of printf's in your code while your debugging your procedures.

Random Numbers. The emulator generates packet loss and errors using a random number generator. Our past experience is that random number generators can vary widely from one machine to another. You may need to modify the random number generation code in the emulator we have supplied you. Our emulation routines have a test to see if the random number generator on your machine will work with our code. If you get an error message:

It is likely that random number generation on your machine is different from what this emulator expects. Please take a look at the routine `jimsrand()` in the emulator code. Sorry.

then you'll know you'll need to look at how random numbers are generated in the routine `jimsrand()`; see the comments in that routine. **Continued on Slide 3-70...**

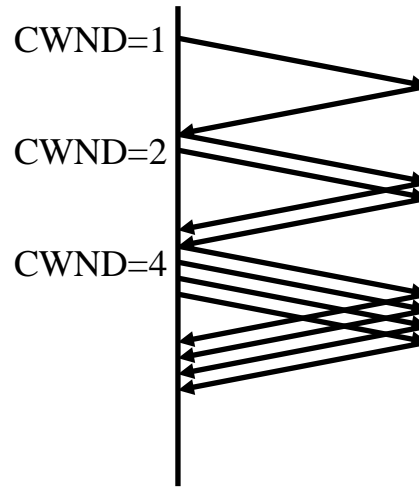
Student Questions

Optional Homework 3D

Try but do not submit.

A TCP entity opens a connection and uses slow start.
Approximately how many round-trip times are required before
TCP can send N segments.

Hint:



Student Questions

- Could you walk through the hint for problem 3D?

Sure.

- Should we discuss it by the case for this question? (i.e., for large or low initial threshold value)

Not sure what the question is.

- Can I get extra credit if I submit it?

No. There is no place to submit it.

Acronyms

- ❑ ACK ACKnowledgement
- ❑ AIMD Additive increase and multiplicative decrease
- ❑ ARQ Automatic Repeat Request
- ❑ CE Congestion Experienced
- ❑ CRC Cyclic Redundancy Check
- ❑ CWND Congestion Window
- ❑ CWR Congestion Window Reduced
- ❑ DA Destination Address
- ❑ DEC Digital Equipment Corporation
- ❑ DECbit DEC's bit based congestion scheme
- ❑ DevRTT Deviation of RTT
- ❑ DNS Domain Name System
- ❑ DP Destination Port
- ❑ ECE Explicit Congestion Experienced
- ❑ ECN Explicit Congestion Notification
- ❑ FIN Final

Student Questions

Acronyms (Cont)

- ❑ FTP File Transfer Protocol
- ❑ GBN Go-Back N
- ❑ HTTP Hyper-Text Transfer Protocol
- ❑ IETF Internet Engineering Task Force
- ❑ IP Internet Protocol
- ❑ ISN Initial Sequence Number
- ❑ kB Kilo-Byte
- ❑ MSS Maximum segment size
- ❑ PBX Private Branch Exchange
- ❑ PSH Push
- ❑ RFC Request for Comments
- ❑ RM Resource Management
- ❑ RST Reset
- ❑ RTT Round-Trip Time
- ❑ SA Source Address
- ❑ SACK Selective Acknowledgement

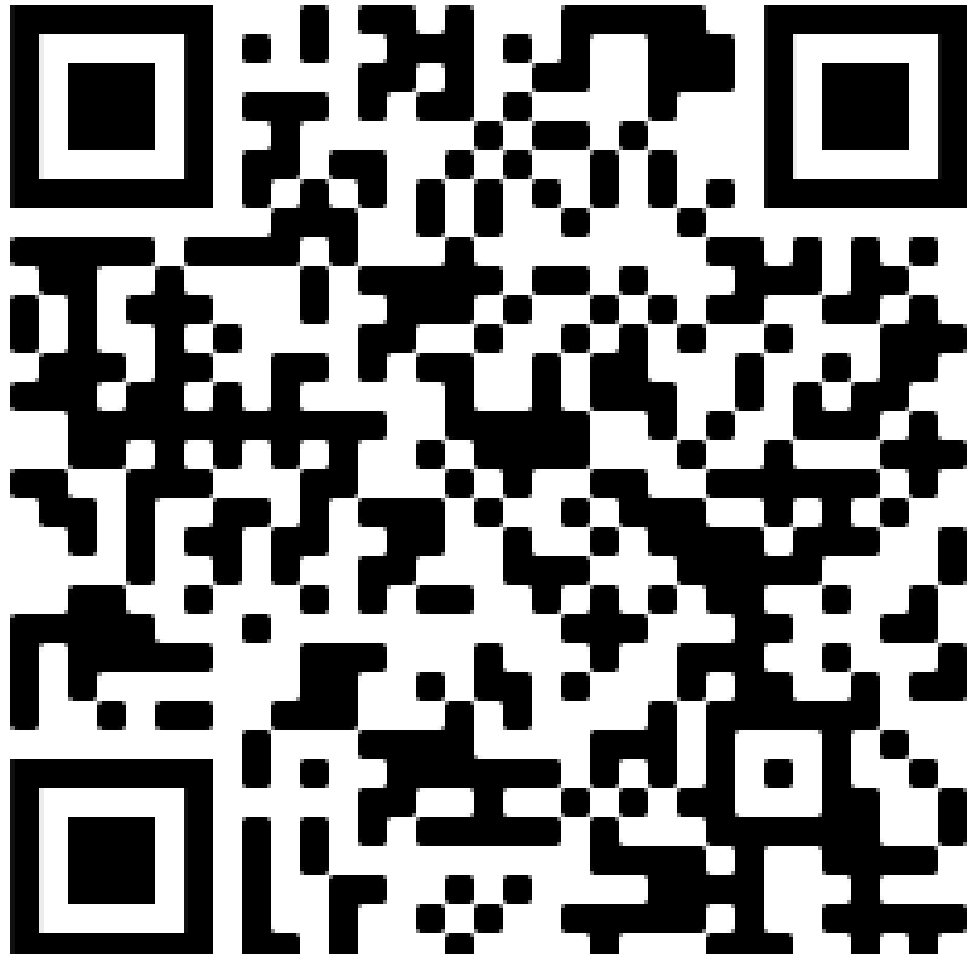
Student Questions

Acronyms (Cont)

- ❑ SMTP Simple Mail Transfer Protocol
- ❑ SP Source Port
- ❑ SStresh Slow Start Threshold
- ❑ SYN Synchronization
- ❑ SYNACK SYN Acknowledgement
- ❑ TCP Transmission Control Protocol
- ❑ UDP User Datagram Protocol
- ❑ URG Urgent
- ❑ VCI Virtual Circuit Identifiers

Student Questions

Scan This to Download These Slides



Raj Jain

<http://rajjain.com>

http://www.cse.wustl.edu/~jain/cse473-22/i_3tcp.htm

Student Questions

- Why neither TCP nor UDP provides bandwidth guarantees service?

TCP/IP community was academic for a long time. Many attempts were made and failed. MPLS has succeeded.

- Domain names are unique across the Internet. What about hostnames? Some people say that in a local area network, hostnames are unique. Is this correct?

Yes. You cannot have two computers with the same name in one LAN. In the Internet, we use "fully qualified domain name (FQDN)." No two computers should have the same FQDN.

- Why does class start at least 5 mins late?
- To achieve quorum and to setup all systems.*

- What is the format of the exam?

See Sample Exam on Canvas.

Related Modules



CSE 567: The Art of Computer Systems Performance Analysis

https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),

https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcg5e_10TiDw



CSE 570: Recent Advances in Networking (Spring 2013)

<https://www.youtube.com/playlist?list=PLjGG94etKypLHyBN8mOgwJLHD2FFIMGq5>

CSE571S: Network Security (Spring 2011),

<https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u>



Video Podcasts of Prof. Raj Jain's Lectures,

<https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw>

Student Questions

Network Utilities

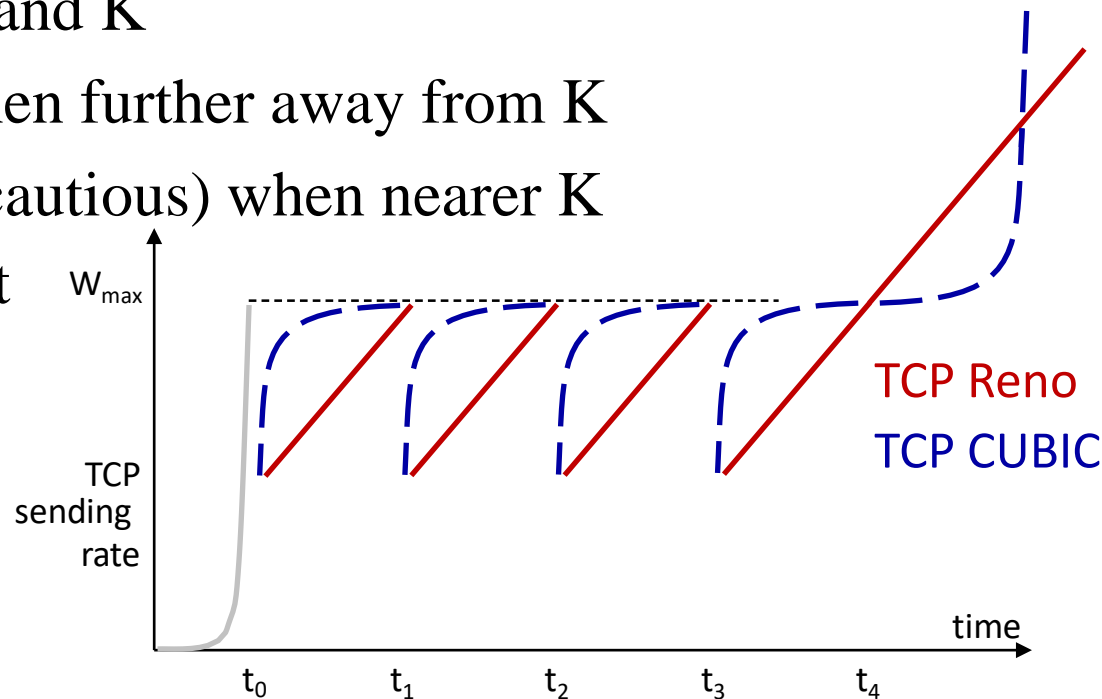
- ❑ TCPview: Shows active ports on your system

<https://docs.microsoft.com/en-us/sysinternals/downloads/tcpview>

Student Questions

TCP CUBIC

- ❑ K : Estimate of time when TCP window size will reach W_{\max} . K is a tunable parameter.
- ❑ Increase W as a function of the *cube* of the distance between the current time and K
- ❑ Larger increases when further away from K
- ❑ Smaller increases (cautious) when nearer K
- ❑ TCP CUBIC default in Linux.
- ❑ Most popular TCP for Web servers



Student Questions

- ❖ How is the estimate of time when TCP window size will reach W_{\max} (K) calculated?

You can compute the time using linear prediction (red line).

Ref: L. Xu, S. Ha, A. Zimmermann, L. Eggert, R. Scheffengger, "CUBIC for Fast Long Distance Networks," RFC 8312, Feb 2018.

Lab 3 Hints

- ❑ Some students received warning messages when trying to compile the C code. This can be fixed by adding the line

```
#include <stdlib.h>
```

and removing all instances of
`char *malloc();`

- ❑ The method `starttimer`, which schedules `timerinterrupt` to trigger, is the one to be careful of.

- `starttimer` schedules `timerinterrupt` to trigger after some amount of time and has the following signature.

- `starttimer(int calling_entity, float increment){ }`

- If `increment` is not cast to a float, it leads to unexpected behavior (triggering back to back to back to back...).

- This is an issue of bit representation. C will take in the bits that represent an int and interpret them as if they were a float. This leads to a passing in a much smaller value than expected.

Student Questions

Lab 3 Hints (Cont)

- i.e., the following will trigger back to back to back...
- `starttimer(A_is_calling_entity, 2000);`
 // Leads to lots of repeated timerinterrupts
- The following will behave as expected
- `starttimer(A_is_calling_entity, (float)2000);` // Casting to float will fix the issue

Student Questions