

SUNetID: _____

CS144
Intro to Computer Networks
Midterm Exam – Wednesday, November 3rd, 2021
Rules: 2 note pages, closed book, computers off

Your Name: _____

SUNet ID: _____@stanford.edu

In accordance with both the letter and the spirit of the Stanford Honor Code, I neither received nor provided any assistance on this exam.

Signature: _____

Check if you would like exam routed back via SCPD: ☐

- The exam has 7 questions totaling 56 points.
- You have 90 minutes to complete them.
- Please keep your answers concise. You may lose points for a correct answer that also includes incorrect or irrelevant information.
- If you would like to make any additional commentary on a multiple-choice answer, please write it below the answer section, but nothing additional is necessary to receive full credit.
- Please box your final answers.

1	/6
2	/6
3	/2
4	/4
5	/8
6	/12
7	/18

Total	/56
--------------	-----

I Layers and abstraction

1. [6 points]:

Which of the following are idempotent operations? Choose **all** that apply.

- A** Making a DNS request to look up the IP address corresponding to a domain name
- B** Sending an HTTP POST request (a client uses a POST request to "post" data to a server, i.e. to submit a form, transfer money from a bank account, write a comment on social media, etc.)
- C** Sending an HTTP DELETE request (a client uses DELETE to delete a specified resource from a server; assume that if a server receives a request to delete a resource that does not exist, it does nothing.)
- D** Inserting a specific element into a set
- E** Incrementing an integer variable by 8
- F** Setting an integer variable to 8

SUNetID: _____

2. [6 points]:

Your computer sends out a DNS request over an Ethernet (link layer) connection.

- (a) What information is included in the first byte of the Ethernet frame's payload?

Circle the best answer.

- A** TCP port
- B** IP version
- C** DNS transaction ID (first field of DNS request)
- D** UDP source port

- (b) What number is in the protocol field of the IP header?

Circle the best answer.

- A** 17 (UDP)
- B** 53 (DNS destination port number)
- C** 1 (ICMP)
- D** 4 (IPv4)
- E** 144 (your favorite class)

- (c) If you were to open Wireshark and examine the packet, what headers would you observe, from outermost to innermost?

3. [2 points]:

You are an engineer at an Internet company, and your team decides to implement an improved, but still reliable, version of TCP. They want this implementation to exist in userspace and to be distinct from the OS-provided TCP functionality.

Coworker A proposes implementing this using UDP sockets (i.e., "TCP over UDP"). This means that the program will open and write TCP segments, including TCP headers, to a UDP socket. The OS will then encapsulate each TCP segment within a UDP datagram before sending it out.

Coworker B argues that this will never work: "TCP and UDP are opposites. TCP is reliable, whereas UDP is unreliable. Our program needs to open a TCP socket if it wants to be reliable."

Who is correct?

A Coworker A

B Coworker B

SUNetID: _____

4. [4 points]:

While watching a recorded lecture on Canvas, you open a second browser tab to scroll through your favorite cat-related blog.

You can assume that these are the only applications connecting to the Internet on your machine, that each tab (Canvas + cat blog) opens TCP sockets, and that your machine has one IP address.

Given this, when a TCP segment arrives, how does the OS figure out which application to deliver it to? (In other words, what does your OS do to make sure that a cat doesn't end up in your lecture video?)

II Routing

In the diagram below, circles represent routers, lines represent links, letters are router names, and numbers indicate link costs.

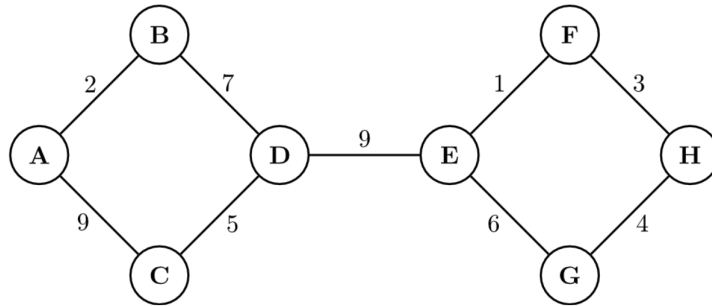


Figure 1: Image description: Graph depicting a network with 8 nodes labeled A through H. Edges connecting nodes as follows: A to B with cost 2, A to C with cost 9, B to D with cost 7, C to D with cost 5, D to E with cost 9, E to F with cost 1, E to G with cost 6, F to H with cost 3, G to H with cost 4.

Throughout the following questions, you may interpret “shortest” and “longest” to mean “lowest-cost” and “highest-cost”.

5. [8 points]:

- (a) For which pair of routers is the shortest path between them the longest?
Please give your answer as a pair of router names; for example ‘A, B’.

SUNetID: _____

- (b) What is the cost of this path (the shortest path between the two routers you identified in part (a))? Please give your answer as a single positive integer.

SUNetID: _____

- (c) Which edges should be removed to create a shortest-path spanning tree rooted at router C?
- A** A–B and E–F
 - B** A–C and E–G
 - C** A–C and D–E
 - D** C–D and G–H
 - E** B–D and G–H
- (d) Imagine this network uses the simplified form of the Bellman–Ford algorithm that we saw in lecture to build its routing tables, in which all routers exchange information with their neighbors in lock-step. If all routers start from scratch (i.e., distance vectors initialized to infinity), how many steps will it take for the network to reach its final routing table configuration? Please give your answer as a single positive integer.

III Fair Queuing

In this question, we introduce you to a simple queuing mechanism that implements fair queuing and has some additional properties. We encourage you to read the entire section before starting the questions.

In this network, all routers support **per-flow round robin (RR)** queuing on the outgoing links and have **sufficient buffer to never cause drops (packet loss)**.

In per-flow round robin (RR), the sending endpoint assigns each packet to a flow, and on the router, each flow has its own queue. The queues are listed in some order. The RR scheduler transmits one packet from the first queue in order, then one packet from the next queue, and so on. After transmitting a packet from the last occupied queue, the RR scheduler repeats the process, starting from the first queue. No queue has priority, and the link rate is shared equally among the queues. **A different queue is maintained for each flow.**

In this scenario, the senders implement no congestion-control protocol. The receiver-advertised window is a fixed amount W —a window big enough to keep the links busy.

Please answer the questions below with respect to the following network diagram:

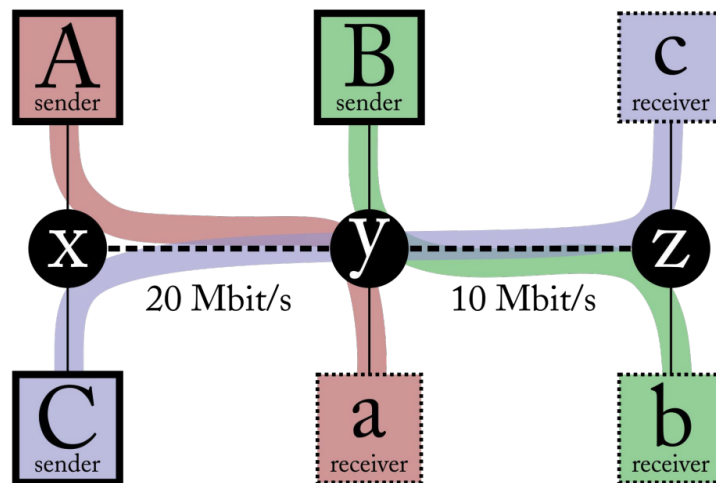


Figure 2: Image description: A network with three flows (A, B, C), three routers (x, y, z), and two links. The flow “A” goes over the link from “x” to “y”, which has a link rate of 20 Mbit/s. The flow “B” goes over the link from “y” to “z”, which has a link rate of 10 Mbit/s. The flow “C” goes over the path from “x” to “y” to “z”, crossing both links.

SUNetID: _____

6. [12 points]:

- (a) What will be the throughput of each flow?

- (b) What is the total throughput of the system?

- (c) What is the goodput (throughput of bytes coming out of the TCP receiver's ByteStream) of each flow and of the system?

- (d) What type of resource allocation or fairness is the outcome of this system?

SUNetID: _____

- (e) Explain why each flow's throughput will converge to this type of fairness

IV Packet Switching and AIMD

In this question, we are going to study a topology with two links and a single switch in between, and a single flow across this topology. We will first calculate the various sources of packet delays across the two links. We will then study the dynamics of the window size curve of a TCP sender that uses AIMD (additive increase, multiplicative decrease).

In the topology below, $r_2 < r_1$. Each link can transmit packets in both directions independently.

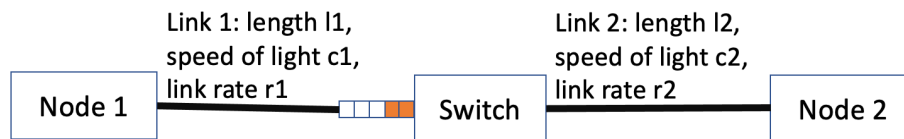


Figure 3: Image description: There are two nodes, node 1 and node 2, with a switch in between (and the switch contains a buffer with a queue for packets that are waiting to be sent out). node 1 is connected to the switch through link 1, and node 2 is connected to the switch through link 2. link 1 has length l_1 and link rate r_1 ; the speed of light through link 1 is c_1 . link 2 has length l_2 and link rate r_2 ; the speed of light through link 2 is c_2 . r_2 is less than r_1 ; each link can transmit packets in both directions independently.

For all answers that require numerical expressions, please express the answer in terms of r_1 , r_2 , l_1 , l_2 , c_1 , c_2 , and any additional variables introduced within the sub-problem.

7. [18 points]:

- (a) What are the components that contribute to the overall delay between when a packet is created by an endpoint and when it is received at another endpoint?

SUNetID: _____

- (b) Assume that the switch is store-and-forward, and assume that there is no other traffic in the network (the router's buffer is always empty). In this case, what is the total, one-way delay for a packet of size p traveling from node 1 to node 2?

- (c) Now assume that there is other traffic in the network, such that the switch's queue always has b packets in it (the switch is still store-and-forward). If the switch buffer constantly has b packets, now, what is the total one-way, end-to-end delay experienced by a packet of size p traveling from node 1 to node 2?

- (d) What is the maximum possible throughput between node 1 and node 2?

- (e) What is the time between when the sender sends a packet, and when it can receive an acknowledgement for the data in that packet, when there is no queuing? This is called the minimum round-trip time (MinRTT).

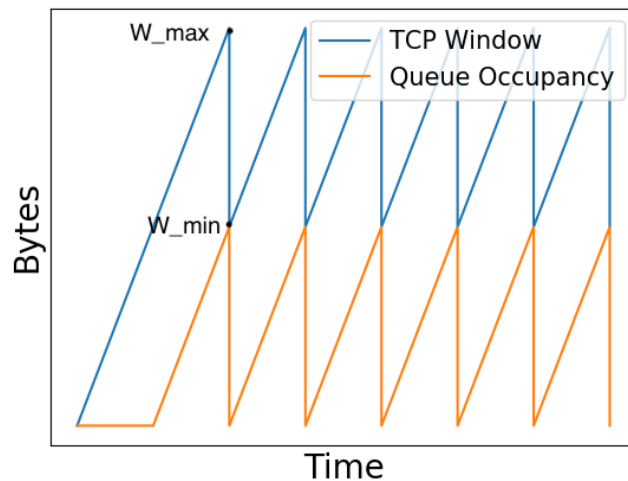


Figure 4: Image description: Picture shows a graph of two lines: the congestion window over time for a TCP AIMD sender, and the queue occupancy of a switch's output queue. The x-axis shows time; the y-axis shows two lines: one for window size and one for queue occupancy (both measured in bytes). At the beginning, the sender's congestion window increases to W_{max} , while the queue occupancy increases to W_{min} . Then, the congestion window drops vertically to W_{min} , while the queue occupancy drops vertically. The graph repeats forever in a sawtooth pattern. The congestion window and queue occupancy begin increasing again until they reach the same thresholds, drop to the same minimum values and then increase to the same maximum values.

- (f) Now let's consider the dynamics of a TCP AIMD sender. Recall that the graph of the congestion window over time looks like a sawtooth (Figure 4). As ACKs are arriving for bytes successfully received, the window increases by 1 segment per RTT. When a loss is observed, the window is reduced. For a standard TCP AIMD sender, what is the relationship between the labeled W_{min} and W_{max} on the curve?

Circle the best answer.

- A $W_{min} = W_{max}/3$
- B $W_{min} = 2 * W_{max}/3$
- C $W_{min} = W_{max}/2$

SUNetID: _____

- (g) What is the optimal buffer size (maximum queue occupancy at the router) such that the link remains always fully utilized, but queueing is minimal? Assume that the delay for an acknowledgement to go from the receiver to the sender is equal to the forward delay without queueing. Please express your answer in terms of the constants given above (e.g. $r1$, $l1$, etc.).

- (h) What happens to W_{min} and W_{max} if $l1$ (the length of link 1) increases (assume that the buffer size at the switch is re-adjusted to a new optimal value, if necessary).

Circle the best answer.

- A** maximum increases, minimum decreases
- B** maximum decreases, minimum increases
- C** both decrease
- D** both increase
- E** nothing

- (i) Why? (two sentences max)