# CS356 Final Exam 2024

**Instructions**
You have 115 minutes to answer the questions.

As announced in Ed Discussion, you may ask us *specific* and *factual* questions during the quiz. If we decide that the question is regarding nitty-gritty details about network protocols, we will answer (e.g. you can ask how a specific routing algorithm works or what port number is used for what protocol).

You may use the one sheet of paper filled with notes that you brought with you. In addition, you may use blank sheets of paper for scrap work. You will not need a calculator for this exam. If you have any questions about the logistics of this quiz, ask us. If you need more space to write your answers, you may attach extra sheets to the question paper.

Please do not cheat. It makes life difficult for everyone. Cheating will be dealt with according to university policies.

All the best!

## Living on the edge?

**Question 1:** Given below are several use-cases. Should we be using an edge deployment or should we use large data-centers?
Assume that both types have the same *total* compute/storage/memory/… capacity, but edge deployments spread them out geographically whereas large data-centers concentrate them in a small number of regions.
Each question is worth 0.5 points. In each case, tick the correct answer.

1. Training a large language model like ChatGPT (Edge/Centralized)
2. A cloud gaming service which renders games on a remote server and streams the video to the user device (Edge/Centralized)
3. Serving a popular video, such as a popular TV show or a live sports broadcast (Edge/Centralized)
4. Serving an unpopular video (such as a lecture by this course's instructor) that has < 50 views spread uniformly over 5 years and viewed from all the different parts of the world (Edge/Centralized)

## Going once… going twice… sold!

**Question 2:** I am trying to bid on spectrum usage rights. I have the following two options. Which one should I value more in terms of the technical capabilities those rights will give me? Note: while in reality, different frequencies have different physical characteristics, we will assume all frequencies are identical. Reminder: Mega is $10^6$ and Giga is $10^9$ (1 point)

(A) 750-950 MHz
(B) 2.7 - 3 GHz

## A question of integrity

**Question 3:** HTML specifies which other objects the browser must load, including scripts and images using “tags”. In the script tag, it is possible to optionally specify an “integrity” attribute that contains a hash function of the script to be loaded. When the browser loads the script, it computes the downloaded script’s hash function and compares it to the one specified in the HTML. If they do not match, it refuses to execute the script. How might this be useful when using a CDN? (1 point)

## Let’s encrypt

A secure, authenticated connection has two properties. 1) Nobody else can understand the content of what they are sending to each other. 2) Nobody else can alter/fake messages and make them look legitimate, i.e., Alice can always tell if the messages she receives are not from Bob (and vice versa for Bob)

The six circumstances on the right are described by two variables: the channel type and the type of pre-shared keys.

There are two types of channels:

1. **No MitM**: others may eavesdrop on conversations, but man-in-the-middle attacks are not possible
2. **MitM**: others may eavesdrop. In addition, they can perform man-in-the-middle attacks

There are three types of pre-shared keys:

1. **None:** Alice and Bob are strangers; Neither party knows any key belonging to the other (public key or a shared symmetric key)
2. **One:** Alice knows Bob’s public key
3. **Both:** Alice and Bob both know each others’ public key

Note: we are not trying to protect meta-data such as when and how much Alice and Bob communicate. We only want to protect *what* they communicate

**Question 4:**
The six rows below ask whether Alice and Bob can establish a secure, authenticated connection in six different circumstances. Write your response as “yes” or “no” in the third column (6 * 0.5 = 3 points)

| Channel type | Pre-shared key type | Is a secure, authenticated connection possible? |
|---|---|---|
| No MitM | None | |
| No MitM | One | |
| No MitM | Both | |
| MitM | None | |
| MitM | One | |
| MitM | Both | |

# BGP strikes back

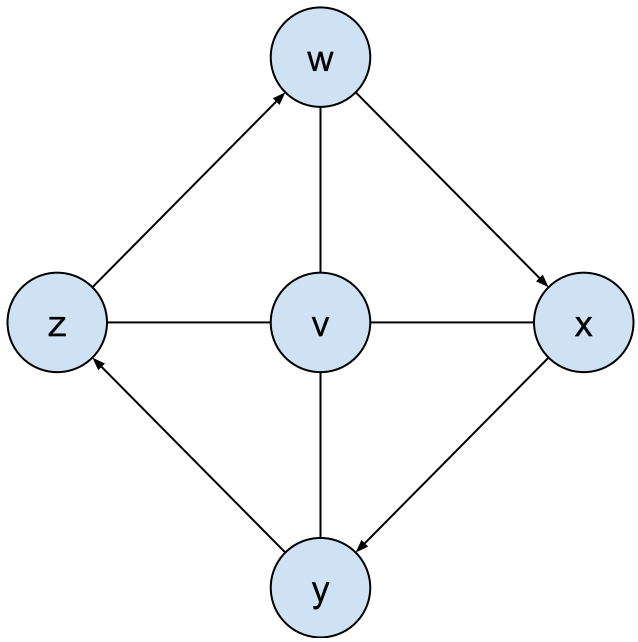


This is the question you were promised would appear in some quiz/final exam. In the above figure, peer-peer relationships are denoted by undirected edges and customer-provider relationships are denoted by a directed edge going from customer *to* provider. Circles represent autonomous systems (ASes).

The questions ask you about whether route flapping can occur when routing to a prefix inside AS v. If you want to answer by giving an example, use the following convention:

For each AS (other than v), write an ordered list of paths that the AS prefers. For instance, if you write for a prefix in v:

w: (w, v), (w, z, v)

It means that w prefers the direct path to v over the path through z. If neither is available, w will not route any packets to that prefix. Note, this example is only for illustrating the notation. It may or may not reflect the AS relationships.

In addition, for each AS, list the neighbors to which it advertises a path to v. Note, your

**Question 5:** Does there exist a set of path preferences and route export policies for the ASes that *can* cause route flapping when routing to a prefix in v? If yes, write one example. If no, argue why not. (points: 1.5 = 0.5 + 1) answers must follow the conventions followed by the two types of relationships (see below)

**Background (covered in class)**:
Customer-Provider relationship

- Customer pays provider for access to the Internet
- Provider exports its customer's routes to everybody
- Customer exports provider's routes only to its customers
- They both export routes to their own prefixes to each other.

Peer-peer relationship:

- Peers exchange traffic between their customers
- An AS never exports providers or other peers' routes to its peer.
- An AS exports a peer's routes *only* to its customers, but isn't obligated to do so

Ranking rule: Everyone prefers customer routes over peer and provider routes. ASes have complete control over their routing preferences as long as they follow this rule.

**Route flapping:** Route flapping occurs when there exists a sequence of route updates such that the routes chosen by the network keeps changing forever.

Here is the content of the above diagram in textual form. The topology has ASes v, w, x y z. AS v is in the center and has a peer relationship with the others. There are directed edges from w to x, from x to y, from y to z and z to w.

**Question 6:** Does there exist a set of path preferences and route export policies for the ASes which guarantees that route flapping can never occur when routing to a prefix in v? If yes, write one example. If no, argue why not. (points: 1.5 = 0.5 + 1)

# To reorder or not to reorder, that is the question

Consider a datacenter that uses a Clos topology as shown below. To maximize load balancing, many datacenters adopt Equal-Cost Multi-Path (ECMP) as a strategy where the packets of each flow are sent through a *randomly chosen* path. In the figure below, there are 3 paths between any two servers. Each time a packet is sent between them, it randomly picks one of the three paths. However, if packets are sent between one of the servers and the WAN (i.e., the internet) there is only one path for them to take

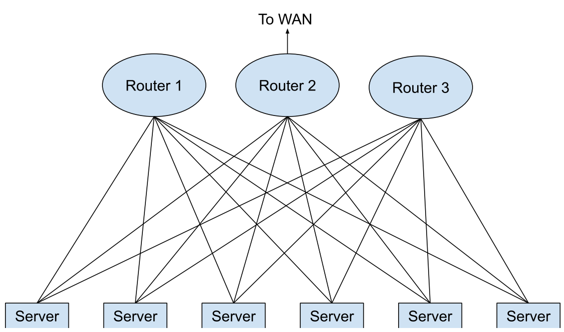


**Assumptions you can make:**

1. TCP uses AIMD and detects loss based on 3 duplicate acknowledgments and timeouts
2. The sending is limited only by the congestion window. The application always has enough data to send.
3. “Full utilization” means that if there is any flow bottlenecked on a given link, that link is fully utilized.
4. There is a TCP flow between every pair of servers, one in each direction. In addition, there are 5 flows from each server to the outside world (WAN)
5. The buffers are large enough that they are not the problem

Hint: There is a reason we ask the first question in this set before the others. Also, the above assumptions are stated for pedantic reasons. You can find the answers without stressing about the specifics.

**Description of figure:** There are 6 servers and 3 routers. Every server is connected to every router. One of the 3 routers is connected to the Wide Area Network (WAN), i.e, the internet

**Question 7:** Suppose there is a network (not necessarily the one shown on the left) where TCP achieves high throughput and ensures fairness between the flows.
Suppose we add a magical box to the path of one of the flows that reorders the packets by shuffling them 10 packets at a time, but leaves everything else unchanged, what will happen? (1 point)

(A) That flow will get more throughput than it was getting before
(B) That flow will get less throughput than it was getting before
(C) The throughput allocation remains exactly (or almost exactly) the same as it as before
(D) All the flows get less throughput than they were getting before
(E) All the flows get more throughput than they were getting before

**Question 8:** What happens to the TCP flows in the network described on the left? (1 point)

(A) All intra-datacenter flows get the same bandwidth and the network is fully utilized
(B) The intra-datacenter flows get different bandwidths, but the network is fully utilized
(C) The WAN flows get almost no bandwidth while the intra-datacenter flows get a disproportionate share of the capacity
(D) Many of the flows get low bandwidth, causing the network to be under-utilized

**Question 9:** Suppose we change the ECMP algorithm so that all packets of one TCP flow go through the same path. However, the path for each TCP flow is randomly chosen. What happens? (1 point)

(A) All intra-datacenter flows get the same bandwidth and the network is fully utilized
(B) The intra-datacenter flows get different bandwidths, but the network is fully utilized
(C) The WAN flows get almost no bandwidth while the intra-datacenter flows get a disproportionate share of the capacity
(D) Many of the flows get low bandwidth, causing the network to be under-utilized

# Cellular networks

Cellular networks use a centralized Medium Access Control (MAC) protocol. Whenever a User Equipment (UE) wants to transmit, it first sends a short packet telling the base station of this need. The base station responds with an allocation of time slots (and frequencies, but we will ignore that detail) that the UE can use.

**Question 10:** Let us define the “uplink latency” to be the time between when a UE wants to transmit a packet to the time when the base station receives it. The “downlink latency” is defined similarly for when the base station wants to transmit to the UE. In an uncongested network, and in the absence of any other information, how would you expect the two to compare? (1 point)

(A) Downlink latency < Uplink latency
(B) Downlink latency = Uplink latency
(C) Downlink latency > Uplink latency

**Question 11:** Suppose there are two base stations that are close enough to each other that some devices can hear both of them. How would you architect the MAC protocol? Any reasonable design suffices (1 point)

## What is the time?

The Network Time Protocol (NTP) is used to synchronize clocks. At a level, a client wanting to synchronize its clock to the real time (i.e. relative to GMT) sends a request to an NTP server. The server responds with the current time t.

The client could set its clock to t. However, that would ignore the network latency. The real time is likely to be t + (time taken between when the server measured its clock and now). As an approximation to this correction factor, the client measures the Round Trip Time (RTT) and sets its clock to t + RTT/2. This assumes that the amount of time taken by the packet to go from the client to the server is the same as the time taken for the return journey.

If we assume that individual links are symmetric (i.e. the latency is the same in both directions), NTP's assumption will hold if the path taken by the packet is the same in both directions. Of course, this depends on the routing algorithm. For the two routing algorithms listed below, answer true/false if it guarantees symmetry after the algorithm has converged (i.e. the routes have stopped changing). Justify your answer by either (informally) proving that it will be symmetric or by giving a counter-example. Each part is worth 1 point for a total of 2 points

**Question 12:** Spanning tree

**Question 13:** BGP

## Did I just lose?

TCP uses two mechanisms to detect packet loss: timeout and duplicate acknowledgments. The following two questions ask what happens if we use only one of the two mechanisms. You should consider questions of both correctness (i.e. will TCP still reliably deliver bytes as long as the network is not completely down?) and performance. You may ignore the impact on congestion control and focus only on the reliability layer

**Question 14:** What happens if TCP uses only timeouts and not duplicate acknowledgments to detect loss? (1 point)

**Question 15:** What happens if TCP uses only duplicate acknowledgments and not timeouts to detect loss? (1 point)

## Dorm room web server

Suppose I am a college student who bought/registered a domain name “something.com”. I want to serve a website on that domain name from a desktop computer kept in my dorm room. Running “ifconfig”, I found that my IP address is 10.10.1.2, which I gave to the Domain Name System (through the service I used to buy the domain name). It complains that this is a *private* IP address and that I am probably behind a NAT.

**Question 16:** Why is it complaining? Why can I not host my website this way? (1 point)

**Question 17:** I go to the website called “whatismyipaddress.com”. It is a service that reads the source IP address from the packets it receives from the HTTP requests, and displays a page containing that source IP along with some other helpful details. It gives me the address 123.162.81.62, which is different from what I saw with ifconfig. What is going on? (1 point)