# CSCI-UA.0480-009 midterm (60 points)

October 17, 2021

## Instructions

1. Write your name and N number on top.
2. Provide concise and clear explanations for all your answers.
3. The questions progress from straightforward to more involved. Pace yourself.
4. Use the other side of each sheet if you need more space.

1. (2 points) What was the primary goal of the Internet when it was designed?

2. (2 points) How were the goals of the Internet different from the goals of the telecommunication network?

3. (2 points) What were some explicit *non-goals* of the Internet?

4. (2 points) If you were building your own private network, what would your goals be and why?

5. (1 point) What is the difference between TCP and UDP?

6. (2 points) Why do we need two transport protocols (TCP and UDP)? Give examples of applications that use TCP and UDP.

7. (5 points) Why does the server end of TCP have 2 sockets associated with it instead of 1 socket like UDP?

8. (5 points) A receiver in the Stop-And-Wait protocol has received the following sequence of packet sequence numbers: 1, 2, 3, 4, 5, 9, 10, 12, 13, 14. Is this a valid sequence? Why or why not? Is it a valid sequence in the sliding window protocol? If the sequence is valid, demonstrate a timeline of network events that leads to this sequence. If the sequence is invalid, explain why the protocol forbids it.

9. Let's say you have a direct 10 Mbit/s (million bits per second) link from your desktop to another desktop. Let's also assume the speed of light is $3 * 10^8$ m/s, and the cable between the two desktops is 5 m long. Answer the following; show your calculations for every question.

   (a) (2 points) What is the propagation delay on this link, i.e., the absolute minimum amount of time that it takes to send a signal from one end of the link to the other?

   (b) (2 points) What is the transmission delay on this link, i.e., the time between when the first and last bits of a packet are sent out on the link? Assume a 1000 bit packet.

   (c) (3 points) What is the queueing delay on this link if (1) you are pinging one desktop from the other with no cross traffic, and (2) if there is cross traffic that creates a queue of size 10 packets, where each packet is 1000 bits?

10. We'll use this and the following question to understand the pros and cons of circuit and packet switching using simple examples. Let's assume we have a single 10 Mbit/s link that can be divided up into 10 logical 1 Mbit/s links. In other words, this link can support a maximum of 10 1-Mbit/s circuits for circuit switching. Let's assume you have several users who want to share this link. First, we'll assume that each of these users wants to send 1 Mbit/s of traffic all the time, i.e., they are never dormant. Now, answer the following questions:

    (a) (1 point) What is the maximum number of users that can be supported by a packet switched network without incurring any queueing delay?

    (b) (1 point) What is the maximum number of users that can be supported by a circuit switched network?

    (c) (1 point) Once the number of users has reached the maximum you just calculated, what happens if one additional user joins the circuit switched network?

    (d) (2 points) Once the number of users has reached the maximum you just calculated, what happens if one additional user joins the packet switched network? Answer this question in two cases: (1) *inelastic traffic*: each user needs exactly 1 Mbit/s and cannot function with anything less, (2) *elastic traffic*: each user can make do with less than 1 Mbit/s.

(e) (1 point) If you had elastic traffic, and you were likely to exceed the maximum number of users, would you use packet or circuit switching? Why?

(f) (1 point) If you had inelastic traffic, and you were likely to exceed the maximum number of users, would you use packet or circuit switching? Why?

11. We'll make the same assumptions as the previous question: 10 Mbit/s link with the ability to support a maximum of 10 1-Mbit/s circuits for circuit switching. But we'll change the traffic profile. Instead of assuming each user is always sending at 1 Mbit/s and is never dormant, we'll assume users are intermittent: each user is on (i.e., the user is transmitting at 1 Mbit/s) for 1 second and then off for 9 seconds, before turning on and off again periodically.

    (a) (1 point) What is the maximum number of such intermittent users that a circuit switched network can support assuming 1-Mbit circuits?

    (b) (3 points) What is the maximum number of such intermittent users that a packet switched network can support without ever building up any queues if the users are perfectly synchronized (Figure 1): every user's on period exactly coincides with every other user's on period?

(c) (4 points) What is the maximum number of such intermittent users that a packet switched network can support without ever building up any queues if the users are desynchronized (Figure 1)? To answer this question, think about how you would play around with the times at which different users are on and off to ensure that the link is never idle.

12. (3 points) Give the precise definition of congestion collapse.

13. (6 points) Give one example of congestion collapse. For this example, specify (1) the individual and aggregate utility functions, (2) how offered load is measured, (3) when and why collapse happens. Sketch out the congestion collapse curve if it helps you explain yourself.

14. Consider the topology of routers in Figure 2. Answer the following questions about how the distance vector protocol would work on this topology. Assume that time starts at 0, and on every time slot or tick, a router can send a distance vector advertisement to its neighbor. This vector is received by the neighbor at the end of that tick, which can then runs its distance vector computation and advertise its own distance vector to its neighbors. These neighbors receive the distance vector at the end of the next tick and so on. To begin with, at time 0, you can assume every router $R$ is initialized with a distance vector with the following entries: (1) if another router $O$ is directly connected to $R$, the distance to $O$ is the link cost of the link between $R$ and $O$, (2) for all other routers, the distance is infinity. Show your calculation at each tick of time.

    (a) (3 points) When does router A learn of some route to D? What is the distance of this route?

(b) (5 points) When does router A learn of the shortest route to D? What is the distance of this route?

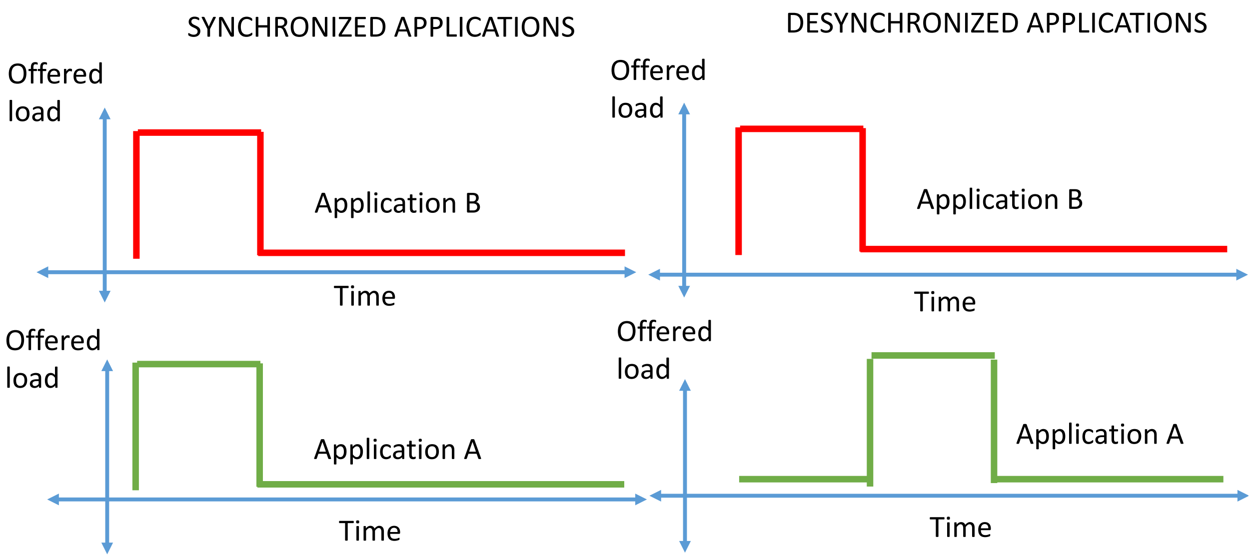


Figure 1: Synchronized and desynchronized applications

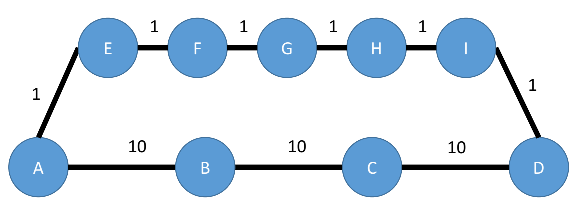


Figure 2: Topology of routers for distance vector question