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

October 23, 2017

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?
   Answer: The primary goals were low-effort interconnection, global connectivity, and generality: the ability to support a wide set of applications.

2. (2 points) How were the goals of the Internet different from the goals of the telecommunication network?
   Answer: Reliability (the telephone network attempted to be available 99.999 % of the time) and performance (it was bad to drop calls). (1 pt for each difference.)

3. (2 points) What were some explicit *non-goals* of the Internet?
   Answer: Security, cost-effectiveness, predictable performance (1pt for 1 non-goal).

4. (2 points) If you were building your own private network, what would your goals be and why?
   Answer: As it is a private network I have more control, so I would like to optimize for performance and cost-effectiveness. If you provide a different answer, but justify it, we will give points for that as well.

5. (1 point) What is the difference between TCP and UDP?
   Answer: TCP: offers reliability
   UDP: doesn't offer reliability (0.5pt each)
   TCP is byte oriented. UDP is packet oriented and preserves packet boundaries.

6. (2 points) Why do we need two transport protocols (TCP and UDP)? Give examples of applications that use TCP and UDP.
   Answer: We need two protocols because of different types of requirements from applications (0.5 pt)
   TCP-Messenger, browsers, email, SSH (information ordering is important)
   UDP- Video Streaming, audio calls, DNS (it is more important to get the most recent frame as opposed to getting every frame.) (0.75 pt each for example)

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

   Answer: In TCP 1 socket is fixed for a server so that it continuosly listens for connection requests and establishes connections to new clients after they have synchronized with the server. Then, the TCP server uses a 2nd data socket to the client through which they communicate after the client and server have synchronized and established a connection. Meanwhile, the original listen socket is still available to listen for connections from new clients. If we only used a single socket then there would be no way for more than one client to connect to a server.

Since each message is independent of every other message in UDP, no synchronization is required before sending the messages, so there is no distinction between the listening socket and the data socket.

(3pt for TCP reasoning and 2pt for UDP reasoning)

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.
   Answer: It is not a valid sequence in the Stop-And-Wait protocol, because the receiver shouldn't accept 9 after receiving 5 by the Stop-And-Wait protocol. Yes it is a valid sequence in sliding window protocol. Assume a window size 14, and the receiver has received only those packets. The remaining might be lost or in transit at the time of recording at receiver, The receiver sends the ACKs back to the sender. Then sender will wait till timeout for retransmissions. Any window size greater than or equal to 5 along with a justification for why this sequence of events happens will be given points. (2pt for stop-and-wait and 3pt for sliding window)

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?
   Answer: (50/3)ns (5 m divided by speed of light in m/s)
   (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.
   Answer: 0.1ms (1000 bits divided by 10 Mbit/s)
   (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?
   Answer: (1) no delay or 0 (2) 1ms (0.1 ms for each packet and 10 packets).

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?
    Answer: 10, because if it's more than 10 the input traffic will exceed 10 Mbit/s, leading to queues building up at the link.
    (b) (1 point) What is the maximum number of users that can be supported by a circuit switched network?
    Answer: 10 because each circuit is 1 Mbit/s and the link itself can carry only 10 Mbit/s.
    (c) (1 point) **Answer:** Once the number of users has reached the maximum you just calculated, what happens if one additional user joins the circuit switched network?
    No service for the additional user. The previous 10 users are unaffected.

(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.
Answer: Inelastic: The service for all the users gets affected. There is 11 Mbit/s of incoming traffic on a 10 Mbit/s. This will lead to continuosly increasing queueing delays and/or packet drops. In either case, this could lead to retransmissions, which would degrade the transport-layer throughput for everyone. If hosts start retransmitting aggrssively, this leads to a congestion collapse.
Elastic: each user will have 10/11 Mbps, assuming the network shares capacity fairly among the users.

(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?
Answer: Packet switching because I want to serve more people. Circuit switching doesn't offer service when users exceed the network's capacity.

(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?
Answer: Circuit switching, because it better to have no service to new users than have no service to all users.

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?
Answer: 10 because each each circuit is designed to handle the worst case traffic requirement, which is 1 Mbit/s for each user.

(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?
Answer: 10 because the network can never support more than 10 Mbit/s of traffic and all senders are on at the same time.

(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.
Answer: 100. During the first second, 10 senders are on. Each sender sends at 1 Mbit/s. During the second, third, ..., tenth seconds, these 10 senders are off. But, during the second second, a new set of 10 senders turn on. Similarly, during the third, a new set of 10 senders turn on, and so on. Because 10 senders turn on at every second, the total number of senders is 100. After the 10th second, this pattern just repeats itself.

12. (3 points) Give the precise definition of congestion collapse.
Answer: Congestion collapse is a term for a situation in which the offered load, i.e., the demand for the networks services, is increasing, but the overall utility of the network is decreasing

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.
Answer: One of the 3 cases described in lecture 6 or anything else that makes sense (too many people talking in a room, too many people watching the NFL superbowl). It doesn't need to be mathematically described as along as it is really clear why there is a congestion collapse. The answer needs to precisely describe what the individual utility functions are (e.g., transport-layer throughput, latency, the ratio of throughput over latency, and so on), what the aggregate utility functions are (e.g., sum of the individual utility functions or max of the individual utility functions), and what the offered load is (total number of sender, total window size).

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?
    Answer: End of 2nd tick. The distance is 30.
    We don't need the contents of all the distance vector tables, so long as you can explain how the distance vectors get updated on each tick clearly.

    (b) (5 points) When does router A learn of the shortest route to D? What is the distance of this route?
    Answer: End of 5th tick. The distance is 6

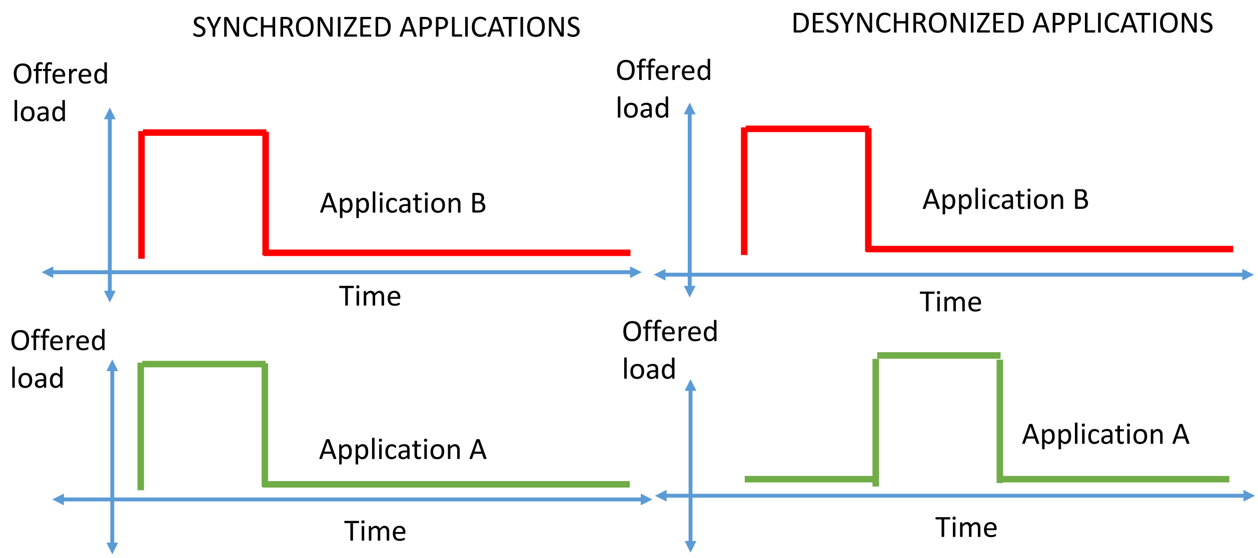

Figure 1: Synchronized and desynchronized applications

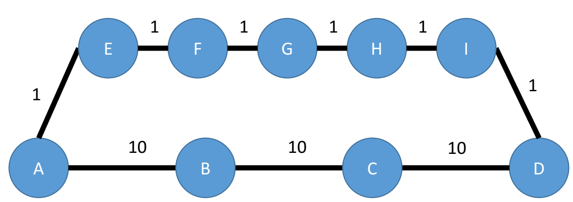

Figure 2: Topology of routers for distance vector question