Study Questions for Exam on Part 1 of ECE 101

Spring 2022
Lectures 1, 7, 8, 9, 10

Table of Contents:

Questions Only

`Q1` [L1] Recall that any information, such as a choice among a group of things, can be represented using a sequence of bits (a binary number). How many bits are needed to represent something from each of the following? *For your convenience, a table of powers of 2 is given:*

2^0	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
1	2	4	8	16	32	64	128	256	512

A. One of the following types of candy bars: Twix, Snickers, KitKat, Almond Joy, Reese's Peanut Butter Cup.

B. One of the following planets: Mercury, Venus, Earth, Mars, Jupiter, Saturn, Uranus, Neptune.

C. One of the days in the year 2022.

D. One of the hours in a day.

`Q2` [L1] Let's say that you want to record just the sound for a 1-hour lecture. Audio is recorded by storing a 1 Byte (8-bit) sample of sound amplitude periodically. For simplicity, let's say that your system samples 50,000 times each second. How many total Bytes do you need to record the lecture? Show your work—*writing the answer as a product is sufficient*.

`Q3` [L7] Denes has some important grade data that he has collected about current UIUC students, and has written a program to analyze those data. Unfortunately, he does not have a powerful enough computer, so he wants to rent a large, powerful computer from Erato's CompuFarm Service.
In one or two sentences (for each answer), explain two separate issues that may complicate the use of Erato's computer by Denes.

`Q4` [L7] Give one reason that using a container to provide a sample of your software product is a better choice than providing the same sample as a virtual machine (VM). Explain your answer in a sentence or two.

`Q5` [L8] In one or two sentences, explain one reason why multicast is difficult to implement in the Internet without using coding theory (the k-out-of-N code approach).

`Q6` [L9] Give **three different examples** that arcs in the social graph might represent, **other than friendship**. Answer with general categories, not specific instances (for example, not "boyfriend").

`Q7` [L9] In one or two sentences, explain why Facebook's TAO architecture breaks the social graph into many more shards than it has computers providing access to the data.

`Q8` [L9] Why is it useful for companies to use computers to identify communities in the social graph when humans already create explicit communities themselves?

`Q9` [L1,L7,L8,L9,L10] For each term below, write the number corresponding to the correct definition from the list below. The definition chosen should correspond to the term's use in lecture and lab. Note that not every definition will be used, as there are more definitions than there are terms.

Terms:

___ A. access control list

___ B. available

___ C. buffering

___ D. client

___ E. clique

___ F. consistent

___ G. datacenter

___ H. graph diameter

___ I. HTTP

___ J. multicast

___ K. pull model

___ L. virtual machine (VM)

Definitions:

1. a group of people who won't talk to anyone outside of the group
2. the protocol used to communicate data, such as web pages, between a web server and a web browser
3. software that abstracts and virtualizes hardware on behalf of applications and protects applications from interfering with one another
4. the maximum possible number of hops required to find a path from any node in a graph to any other
5. an approach in which updates to information are only delivered to a client on request, rather than being proactively sent as soon as changes are made
6. a co-located group of computers and storage used to provide Internet service to many users simultaneously
7. a property provided by Internet storage services, guaranteeing that stored data are not corrupted
8. a software program that emulates a computer and associated hardware, allowing execution of untrusted software without endangering the physical system or any data stored on it
9. sending information from one source to more than one recipient
10. the minimum bandwidth available on any edge in a graph
11. a property provided by Internet storage services, guaranteeing that stored data are

accessible at any time and from any place

12. a list of rules used to determine whether a user is allowed to look at and possibly modify a given document or service
13. Harry Truman Tercius Potter, a fictional character created by J.K. Rowling
14. a property provided partially by Internet services, ensuring that all users see the same version of data at the same time
15. a model in which two groups of college students suspend a rope across a mud puddle and try to pull the other group into it
16. a group of nodes in a graph in which every pair of nodes in the group is connected by an edge in the graph
17. a technique used to stream video or other forms of content in the Internet, in which a small delay is used to smooth out variable needs and bandwidth
18. a computer that makes use of a service provided by another computer over the Internet

`Q10` [L8] Consider the picture shown below, in which three pairs of users are trying to send videos simultaneously through a part of the Internet.

- Jan wants to send a 200 MB video to Pat (along the dotted black line),
- Alice wants to send a 200 MB video to Bob (along the solid, light grey line), and
- Erato wants to send a 200 MB video to Denes (along the double, dashed, dark grey line).

Their routes share links as shown.

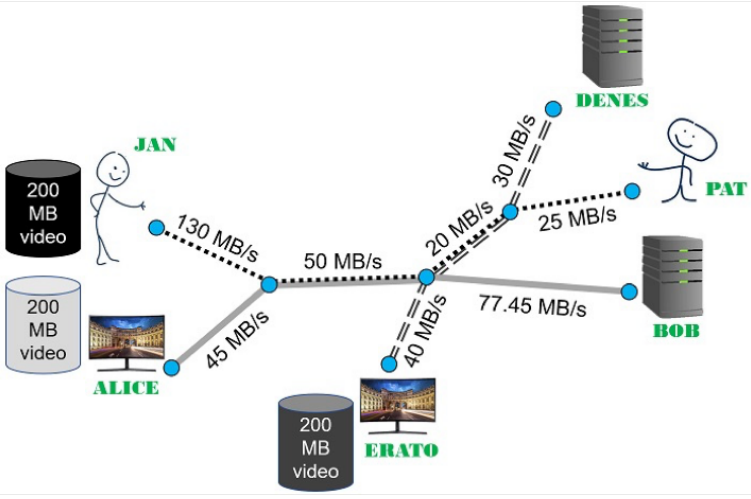

Assume for all questions that TCP is able to reach the rates shown on the links, and that link sharing is fair unless one of the connections is limited by another link along its route.

A. At what rate can Jan send to Pat?

B. At what rate can Alice send to Bob?

C. At what rate can Erato send to Denes?

D. How long does it take Alice to send the full 200 MB video to Bob?

`Q11` [L10] In some cases, a disk may produce a random set of bits when asked to return previously stored data. In a sentence or two, explain how a global Internet service provider can guarantee the reliability of a client's data despite such failures.

`Q12` [L8] You have received 5 of 7 bits encoded using the variant of the Hamming code discussed in lecture, in which each circle should contain an even number of 1 bits. Identify the missing bits so that you can recover the original message, then write the 4-bit message below. Show your work (write bits into the circles as well as on the lines).

Message: __ __ __ __

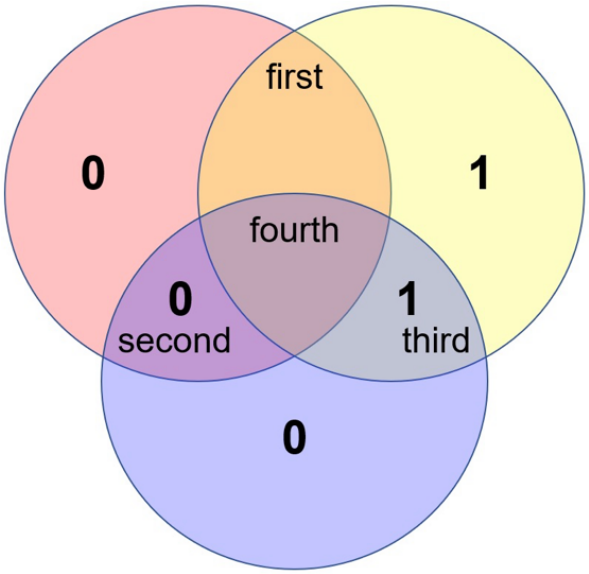

Questions + Solutions

`Q1` [L1] Recall that any information, such as a choice among a group of things, can be represented using a sequence of bits (a binary number). How many bits are needed to represent something from each of the following? *For your convenience, a table of powers of 2 is given:*

2^0	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
1	2	4	8	16	32	64	128	256	512

A. One of the following types of candy bars: Twix, Snickers, KitKat, Almond Joy, Reese's Peanut Butter Cup.

> 5 brands, so 3 bits.

B. One of the following planets: Mercury, Venus, Earth, Mars, Jupiter, Saturn, Uranus, Neptune.

> 8 choices, so 3 bits.

C. One of the days in the year 2022.

> 365 days, so 9 bits.

D. One of the hours in a day.

> 24 hours, so 5 bits.

`Q2` [L1] Let's say that you want to record just the sound for a 1-hour lecture. Audio is recorded by storing a 1 Byte (8-bit) sample of sound amplitude periodically. For simplicity, let's say that your system samples 50,000 times each second. How many total Bytes do you need to record the lecture? Show your work—*writing the answer as a product is sufficient*.

> 1 B × 50000 samples / second × 60 minutes / hour × 60 seconds / minutes
>
> (180,000,000 Bytes = 172 MB; for the raw audio channel; audio, as with most recordings, is usually compressed; modern audio typically has several parallel channels)

`Q3` [L7] Denes has some important grade data that he has collected about current UIUC students, and has written a program to analyze those data. Unfortunately, he does not have a powerful enough computer, so he wants to rent a large, powerful computer from Erato's CompuFarm Service.
In one or two sentences (for each answer), explain two separate issues that may complicate the use of Erato's computer by Denes.

- Denes' data are sensitive! Denes must trust Erato with private information about the UIUC students.
- Denes' program may be valuable. Denes must trust Erato not to steal the program (keep a copy) for Erato's own purposes.
- Erato's computer may have different hardware, OS, and/or software from Denes'. It's not clear that Denes' program will work when run on Eratos' machine.

`Q4` [L7] Give one reason that using a container to provide a sample of your software product is a better choice than providing the same sample as a virtual machine (VM). Explain your answer in a sentence or two.

- VMs are much larger than containers, so downloading a VM is more of a hassle for the customers.
- In order to provide a VM, I need to have the right to ship both the OS and all other software contained within the VM to my customers, not just the right to ship my own software.

`Q5` [L8] In one or two sentences, explain one reason why multicast is difficult to implement in the Internet without using coding theory (the k-out-of-N code approach).

- Without support, multicast using multiple transmissions has to share the first link's (and maybe more) bandwidth amongst all users.
- Multicast has difficulty with lost packets, since every user may lose a different set.
- Multicast has difficulty managing the variations in bandwidth along the routes to each recipient.
- Multicast implementations often rely on routers in the Internet performing special actions to support the multicast, but Internet routers see millions of connections and cannot handle individual connections in special ways.

`Q6` [L9] Give **three different examples** that arcs in the social graph might represent, **other than friendship**. Answer with general categories, not specific instances (for example, not "boyfriend").

> - Marriage/relationship
> - Family relationship (mother/father/son/daughter/etc.)
> - Worked for
> - Founded
> - Studied at
> - Taught at
> - (lots of others—should not be hard to get full points)

`Q7` [L9] In one or two sentences, explain why Facebook's TAO architecture breaks the social graph into many more shards than it has computers providing access to the data.

> Shards have variable size, so using only one shard per computer leads to unnecessarily heavy load on some computers, which can (and is) fixed by using many shards per computer and balancing the load.

`Q8` [L9] Why is it useful for companies to use computers to identify communities in the social graph when humans already create explicit communities themselves?

> Humans do not identify all communities explicitly, and may not even realize that they are a part of a community in some cases. These communities may still be of interest to the humans, or may have commercial value to the companies.

`Q9` [L1,L7,L8,L9,L10] For each term below, write the number corresponding to the correct definition from the list below. The definition chosen should correspond to the term's use in lecture and lab. Note that not every definition will be used, as there are more definitions than there are terms.

Terms:

___ A. access control list

___ B. available

___ C. buffering

___ D. client

___ E. clique

___ F. consistent

___ G. datacenter

___ H. graph diameter

___ I. HTTP

___ J. multicast

___ K. pull model

___ L. virtual machine (VM)

Definitions:

1. a group of people who won't talk to anyone outside of the group
2. the protocol used to communicate data, such as web pages, between a web server and a web browser
3. software that abstracts and virtualizes hardware on behalf of applications and protects applications from interfering with one another
4. the maximum possible number of hops required to find a path from any node in a graph to any other
5. an approach in which updates to information are only delivered to a client on request, rather than being proactively sent as soon as changes are made
6. a co-located group of computers and storage used to provide Internet service to many users simultaneously
7. a property provided by Internet storage services, guaranteeing that stored data are not corrupted
8. a software program that emulates a computer and associated hardware, allowing execution of untrusted software without endangering the physical system or any data stored on it
9. sending information from one source to more than one recipient
10. the minimum bandwidth available on any edge in a graph
11. a property provided by Internet storage services, guaranteeing that stored data are accessible at any time and from any place
12. a list of rules used to determine whether a user is allowed to look at and possibly modify a given document or service

13. Harry Truman Tercius Potter, a fictional character created by J.K. Rowling
14. a property provided partially by Internet services, ensuring that all users see the same version of data at the same time
15. a model in which two groups of college students suspend a rope across a mud puddle and try to pull the other group into it
16. a group of nodes in a graph in which every pair of nodes in the group is connected by an edge in the graph
17. a technique used to stream video or other forms of content in the Internet, in which a small delay is used to smooth out variable needs and bandwidth
18. a computer that makes use of a service provided by another computer over the Internet

> *12* A. access control list
>
> *11* B. available
>
> *17* C. buffering
>
> *18* D. client
>
> *16* E. clique
>
> *14* F. consistent
>
> *6* G. datacenter
>
> *4* H. graph diameter
>
> *2* I. HTTP
>
> *9* J. multicast
>
> *5* K. pull model
>
> *8* L. virtual machine (VM)

`Q10` [L8] Consider the picture shown below, in which three pairs of users are trying to send videos simultaneously through a part of the Internet.

- Jan wants to send a 200 MB video to Pat (along the dotted black line),
- Alice wants to send a 200 MB video to Bob (along the solid, light grey line), and
- Erato wants to send a 200 MB video to Denes (along the double, dashed, dark grey line).

Their routes share links as shown.

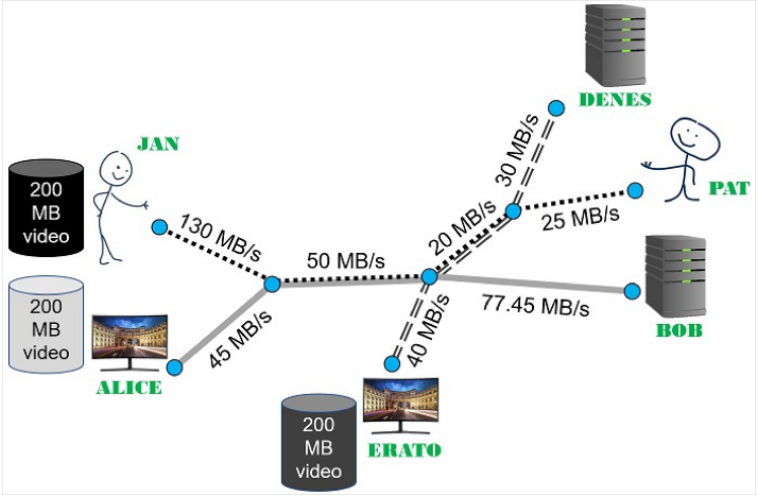

Assume for all questions that TCP is able to reach the rates shown on the links, and that link sharing is fair unless one of the connections is limited by another link along its route.

A. At what rate can Jan send to Pat?

B. At what rate can Alice send to Bob?

C. At what rate can Erato send to Denes?

D. How long does it take Alice to send the full 200 MB video to Bob?

> Looking along their routes, Jan and Erato are both limited by the 20 MB/s link. Alice is limited to 45 MB/s by the first link along the route, but sharing the 50 MB/s link with Jan is likely to impose a lower limit. Since Jan and Erato share fairly, each is limited to only 10 MB/s. Thus, on the 50 MB/s link shared by Jan and Alice, Alice is able to transmit 40 MB/s, which limits the overall route, as it is below the first link's limit of 45 MB/s.
>
> A. 10 MB/s
>
> B. 40 MB/s
>
> C. 10 MB/s
>
> D. 200 MB / 40 MB/s = 5 seconds

`Q11` [L10] In some cases, a disk may produce a random set of bits when asked to return previously stored data. In a sentence or two, explain how a global Internet service provider can guarantee the reliability of a client's data despite such failures.

> Data are spread across multiple disks using a k-out-of-N code (for example), so a single disk does not result in data corruption.

`Q12` [L8] You have received 5 of 7 bits encoded using the variant of the Hamming code discussed in lecture, in which each circle should contain an even number of 1 bits. Identify the missing bits so that you can recover the original message, then write the 4-bit message below. Show your work (write bits into the circles as well as on the lines).

Message: __ __ __ __

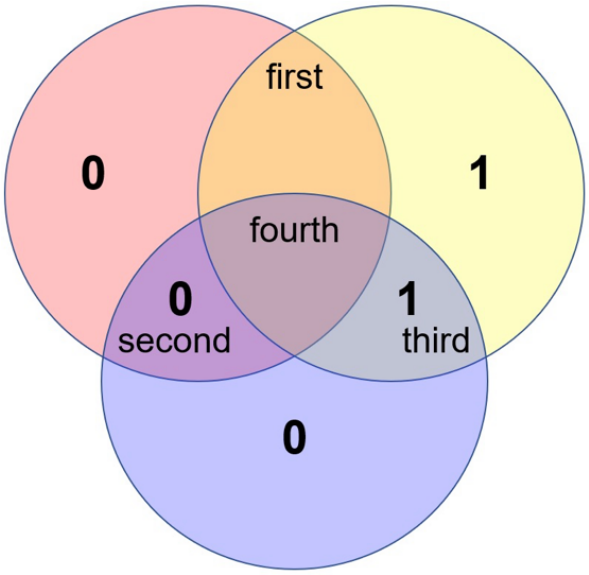

> The “fourth” bit of the message must be 1 in order to have two 1 bits in the blue circle. Both the red and the yellow circle then require that the “first” bit be 1 as well. Message is 1 0 1 1.