Stimulus

The Physical Layer

Recap

When transmitting bits over a physical medium, the transmitted signal can be characterized by the frequencies it occupies. Signals that are sent over a disjoint set of frequencies do not interfere with each other. If a receiver gets two signals sent at the same frequency and at the same time, it usually cannot decode either signal.

Suppose a signal only contains frequencies between f_1 and f_2, and does not contain any frequencies outside this range, then its bandwidth is said to be $B = f_2 - f_1$. If the signal strength is S and noise is N, then the maximum number of bits per second it can communicate is given by

$B \log_2 (S / N)$

When transmitting over the air, there is usually a practical limit to how large the frequency can be. Most medium to long distance links require that the signal not contain frequencies above 10 GHz.

Note: the reason a receiver that receives two signals cannot decode both of them is that the SNR becomes S / (N + signal from the other transmission). If the signal from the other transmission is large, SNR becomes unacceptably low.

Scenario 1:

Suppose there are four radio devices A, B, C and D. They have the right to transmit in a fixed bandwidth range.

(A, B) and (C, D) are close to each other (say, within 100 ft). However A and B are far away from C and D (say 1000 ft). Further, between them is a wall that reduces signal strength by 10x. This is represented by the ASCII art below

```
A       |       C
B       |       D
```

They are limited to transmitting at a maximum power which makes the maximum SNR between neighboring pairs to 1000 and therefore between the distant pairs to 1. This is because signal strength decreases as 1 / distance^2.

1 Multiple Answer 2 points

SpaceX is a space launch company that has deployed a large number of low earth orbit (LEO) satellites to offer internet services to a large portion of the earth. It has approached several governments to ask for the right to transmit at certain frequencies. What *technical* considerations should the governments make? Note: the question is only asking about the *technical* considerations and not for opinions on what is right/advantageous/...

For concreteness' sake, suppose it is asking the US government for the right to transmit between 2.4 GHz and 2.42 GHz and they have 7000 satellites deployed all over the earth and the government can assume that this will not increase beyond 34000.

- [ ] They should be generous with granting SpaceX the right since it does not cost anything, and it could offer connectivity to all rural areas
- [ ] They should be stingy with granting the right since transmissions from space will interfere with all cellular transmissions on earth
- [x] They should be careful with granting the right since spectrum is a scarce resource that needs to be allocated carefully
- [ ] Internet from space will offer a serious competitive threat to incumbent telecom operators in ***urban*** areas since users will want to use SpaceX as their primary internet provider
- [x] Internet from space will offer a serious competitive threat to incumbent telecom operators in ***rural*** areas since users may want to use SpaceX as their primary internet provider

2 Multiple Answer 1 point

Consider scenario 1 described on the left. Suppose A and B want to talk to each other and C and D want to talk to each other. How do we *maximize* the total bandwidth between the pairs?

Select all options that are good. Select multiple if they are equally (or almost equally) good. Do not worry about the overheads of correctly allocating frequency/time

- [ ] Give half of the frequency to each pair so they can communicate simultaneously without interference
- [ ] Allocate half of the time to each pair so that they can communicate using all the frequencies
- [x] Both pairs can communicate at all the frequencies the entire time since interference will be minimal
- [ ] Split both the frequency and time in half to maximize flexibility

3 Multiple Answer 1 point

Consider scenario 1 described on the left. Suppose A and C want to talk to each other and B and C want to talk to each other. How do we *maximize* the total bandwidth between the pairs?

Select all options that are good. Select multiple if they are equally (or almost equally) good. Do not worry about the overheads of correctly allocating frequency/time

- [x] Give half of the frequencies to each pair so that they can communicate simultaneously without interference
- [x] Allocate half of the time to each pair so that they can communicate using all the frequencies
- [ ] Both pairs can communicate at all the frequencies the entire time since interference will be minimal
- [x] Split both the frequency and time in half to minimize interference

4 True or False 0.5 points

Error correction requires fewer redundant bits than error detection.

- ( ) True
- (•) False

5 True or False 0.5 points

To detect *all* 1 bit errors in a file of N bits, we need O(N) redundant bits

- ( ) True
- (•) False

6 True or False 0.5 points

It is possible to ***detect*** nearly all errors in a 1 GB file using just 32 bits of redundant code

- (selected) True
- False

7 True or False 0.5 points

It is possible to ***correct*** nearly all errors in a 1 GB file using just 32 bits of redundant code

- True
- (selected) False

8 True or False 1 point

A MAC protocol that, after collision, waits with probability 0.5 before transmitting again in every subsequent slot will not work well when the number of transmitters are much larger than 3

- (selected) True
- False

9 True or False 1 point

If each transmitter is given a unique frequency on which it can transmit and nobody else is allowed to, then no MAC protocol is needed

- (selected) True
- False

Stimulus

Congestion Control

Consider a link where the Round Trip Time (RTT) is 40 ms when the queue is empty and the bottleneck link rate is 24 Mbit/s. The buffer size is 80 Kbytes.

There is only one active flow on this entire network path

Recap:

The AIMD algorithm we discussed in class was the following

Initialize cwnd = 1500 bytes # 1 packet

def on_packet_acked():
  if it has been more than 1 RTT since the last cwnd update and no packets have been detected to have been lost from this acknowledgment:
    cwnd = cwnd + 1

def on_packet_loss_detected_by_dupack():
  if the last cwnd decrease was over 1 RTT ago:
    cwnd = cwnd / 2

def on_timeout():

  cwnd = 1500 bytes

10 Numeric 1 point

If there is only a single flow on this link, we are guaranteed to fully utilize the link if the congestion window is greater than X. What is X in bytes?

120,000

11 Numeric 1 point

We are guaranteed to not lose any packets as loss as the congestion window is less than Y. What is Y in bytes?

200,000

12 True or False 1 point

If the sender were to run AIMD on this link it will fully utilize the link

- ○ True
- ● False

Stimulus

Reliable delivery

Suppose X is sending data to Y over a reliable byte-stream like TCP or QUIC. The questions on the right specify what acknowledgments (ACKs) X receives from Y.

In both cases, describe (a) what X can conclude and (b) what it assumes about the network to reach this conclusion

Note, your answer to (a) can depend on what assumptions you are willing to make in (b).

13 Essay 1 point

X receives the following two ACKs from Y that tells it what sequence numbers have been received:

First ACK:
[0, 10000], [12000, 11000]

Second ACK:
[0, 10000], [12000, 13000]

5 min / 200 max (word limit)
Must use at least 5 words

14 Essay 1 point

After the first two ACKs, it receives a third ACK as follows:
[0,10000], [12000, 13000], [14000-15000]

5 min / 2000 max (word limit)
Must use at least 5 words

15 Multiple Choice 1 point

It is possible to send two HTTP requests back-to-back on the same connection. When this happens, the server will simply respond with two HTTP responses back-to-back on the same TCP connection

For example, suppose I open a TCP connection to example.com at port 80 and send the following text:

```
GET / HTTP/1.1

Host: example.com

User-Agent: CustomClient/1.0

Accept: */*

GET /image.png HTTP/1.1

Host: example.com

User-Agent: CustomClient/1.0

Accept: */*
```

The server will first send the response for the home page and then it will send the image. Now, suppose that one of the packets containing data about the first response gets lost, but all other packets are delivered without loss. How will TCP interface with the higher layer behave?

- (●) It will wait for that packet to be retransmitted and received before delivering the second request to the application
- ( ) It will deliver the second request to the application immediately before the lost packet is retransmitted since the second request is intact
- ( ) It will send an HTTP error message to the application saying that the requested object could not be received
- ( ) The TCP connection will break, which tells the application to make another connection

Stimulus

Routing symmetry

Consider two network participants X and Y. The questions on the right ask whether the path from X to Y will be the same as the path from Y to X for two different routing algorithms

16 True or False 0.5 points

Spanning tree routing

- (selected) True
- False

17 True or False 0.5 points

BGP

- True
- (selected) False