CS118 Midterm

TOTAL POINTS

94 / 120

QUESTION 1

1 20 pts

1.1 1.1 (TCP as a transport) 3 / 3

- 0 **correct**
- 0.5 Incorrectly select
- 0.5 Incorrectly select
- 0.5 Incorrectly select
- 1 should select 3 correct options
- 1 should select 3 options
- 1 should select 3 options

1.2 1.2 (stateful) 0 / 3

- 0 Correct
- 0.5 Select a wrong answer
- 0.5 Select a wrong answer
- 0.5 Select a wrong answer
- 3 **Wrong answer, leave blank**

1.3 1.3 (p2p) 1 / 1

- 0 **Correct**
- 1 Incorrect (b is not chosen)
- 0.5 At least one other (incorrect) options selected

1.4 1.4 (HTTP request/response info) 2 / 3

- 0 Correct
- 1 missing d
- 1 missing e
- 0 **missing f**
- 1 **One wrong answer (a-c)**
- 2 Two or more wrong answers (a-c)

1.5 Protocol is ... 2 / 2

- 1 At least one correct (format/order/actions)
- 0 **Two or more correct (format/order/actions)**
- 2 Incorrect

1.6 Most common HTTP method 1 / 2

- + 1 **GET**
- + 1 POST
- + 0 Incorrect/Blank

1.7 DNS and HTTP 1 / 2

- + 0.5 replicated
- + 0.5 **cached**
- + 0.5 replicated
- + 0.5 **cached**
- + 0 Incorrect

1.8 Common TCP/UDP functions, unique TCP functions 4 / 4

- + 2 **multiplexing/demultiplexing/error detection**
- + 1 half credit to common function
- + 1 **delivery guarantee**
- + 1 **flow control**
- + 1 **congestion control**
- + 0 Incorrect
- + 1 Incorrect: send/connect/accept/listen calls

QUESTION 2

2 20 pts

2.1 2.1 Delay for 12th packet 3 / 6

- + 6 Correct
- + 3 **Correct transmission delay**
- - 1 Incorrect number of RTTs
- + 0 Incorrect

2.2 2.2 Total delay 2 / 2

- + 2 **Correct 501ms, accepted answer if +-100ms**
- + 0.75 Correct transmission delay
- + 0.75 Correct total number of packets
- + 0 Math problem
- + 0 Incorrect

2.3 2.3 Total delay with 100ms propagation delay 6 / 6

- 0 **Correct (correct 2.64s, accepted 2-3s)**
- 1 No or incorrect explanation provided

- 3 Wrong value, but within reasonable range from the correct (1-2, 3-5)
- 0.5 Incorrect RTT calculation
- 6 Incorrect

2.4 2.4 Shortest delay / adjust window 3 / 6

- 0 Correct (20-21, delay ~600ms)
- 0.5 No (or incorrect) optimal delay calculated
- 1.5 Mentioned value >21, but no or incorrect explanation provided
- 6 Incorrect
- **3 Mentioned to increase window, but <21 or way too many**

QUESTION 3

3 20 pts

3.1 3.1 Query for amazon.com/A 4 / 4

- **0 Correct**
- 2 Second query incorrect
- 1 Problem with one of the queries
- 4 Incorrect / no answer

3.2 3.2 Query for google.com/MX 3 / 3

- **0 Correct**
- 1 Issue with the answer (one unnecessary query)
- 1.5 Didn't include query t google.com NS
- 3 Incorrect / missing / more than one unnecessary query

3.3 3.3 Query for mail/hangout.google.com/AAAA 5 / 5

- **0 Correct**
- 1 Extra (or missing) query for 1st query
- 1 Extra (or missing) query for 2nd query
- 2 No more than two unnecessary queries for one of the queries
- 3 More than 2 unnecessary queries for one
- 5 Incorrect / missing

3.4 3.4 List cached records 4 / 5

- **+ 2 google.com/MX, primary.google.com/A, backup.google.com/A (from 3.2)**
- **+ 1 google.com/NS**
- **+ 2 mail.google.com/AAAA, hangout.google.com/AAAA**
- 0.75 One wrong domain
- 1.5 Two wrong domains
- **1 Type of records not specified**
- + 0 Incorrect / missing

3.5 3.5 Reachable 3 / 3

- **0 Correct**

QUESTION 4

4 20 pts

4.1 4.1 Sequence numbers 10 / 10

- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- **+ 1 One correct sequence number of flag**
- + 0 Incorrect/blank

4.2 4.2 Missing exchange 3.5 / 5

- 0 Correct
- 0.5 One or more incorrect / unnecessary / missing exchanges
- 0.75 One flag or sequence number is wrong
- **1.5 More than one flag or sequence number is wrong**
- 2.5 Sequence numbers / flags (or their relation) not shown
- 5 Incorrect / missing

4.3 4.3 Max TCP pipeline 3 / 5

- 0 Correct (or close)
- 2 Too large or too small
- 2 Incorrect statement that there is no limit
- **2 The result is not throughput (bit/s or byte/s)**
- 0.5 Right direction, but didn't give complete answer
- 5 No attempt
- 3 Attempted, but didn't give the answer

QUESTION 5

5 20 pts

5.1 5.1 UDP checksum 4 / 5

+ 1.5 **Found checksum in the packet**

+ 4 Attempted to calculate 1-complement of 1-complement sum of 16-bits

+ 1 Gave an answer without basis

+ 2.5 **Calculated sum, but didn't indicate 1-complement of 1-complement / not correct items added**

+ 2.5 Attempted calculate but not 1-complement of 1-complement sum or not 16-bit numbers

+ 0 No attempt / incorrect

5.2 5.2 IPv4 header checksum 5 / 5

+ 1.5 **Found checksum in the packet**

+ 4 **Attempted to calculate 1-complement of 1-complement sum of 16-bits**

+ 1 Gave an answer without basis

+ 2.5 Calculated sum, but didn't indicate 1-complement of 1-complement / not correct items added

+ 2.5 Attempted calculate but not 1-complement of 1-complement sum or not 16-bit numbers

+ 2 Give semi-valid answer not to the question asked

+ 0 No attempt / Incorrect

5.3 5.3 Demultiplex 6 / 6

- 0 **Correct**

- 0 **Not mentioned that OS uses UDP-specific lookup table to find app socket**

- 3 Incorrect mentioning of sourceIP/sourcePort as part of the lookup

- 1 Didn't mention destination IP for demultiplexing

- 3 Didn't mention use of destination ip&port to lookup in kernel's UDP socket-app table

- 4 Only showed port number

- 6 Incorrect / missing

5.4 5.4 UDP facts 1 / 4

+ 1 UDP payload (2^16)

+ 1 **Port numbers (0, 2^16-1), ok if start 1024**

+ 1 Number of distinct apps (2*2^16)

+ 1 Number of apps to prevent (2^16)

+ 0 Incorrect

QUESTION 6

6 20 pts

6.1 6.1 Secret message 3.5 / 4

- 0 Correct

- 1 Mentioned public key encryption, but didn't discuss how to get public key for a person you never met before

- 0.5 **Mentioned how to get public key, but didn't define why it should be trusted**

- 4 Incorrect / missing

- 3 Mention PGP, but in incorrect context

6.2 6.2 Info from email 2 / 4

+ 4 Correct

+ 2 **Mentioned PGP signing and telling (out-of-band) your key (key fingerprint) to the professor / out-of-band acknowledging sending email**

+ 2 At least two objective info items listed

+ 1 At least one objective info listed

+ 0 Incorrect / missing

6.3 6.3 Invalid HTTPS 2 / 4

- 0 Correct (At least 2 reasons listed)

- 2 **Only one correct reason listed**

- 4 Invalid / missing answer

6.4 6.4 Multiple DNS records for youtube.com/A 3 / 4

- 0 Correct (at least 3 correct reasons listed)

- 1 **Only two correct reasons listed**

- 2 Only one correct reason listed

- 0 Incorrect/missing answer

6.5 6.5 HTTP/2 vs QUIC 4 / 4

- 0 **Correct**

- 1.5 Only one reason listed

- 4 Incorrect/missing answer

CS118
Spring 2017 Midterm Exam

1 hour 50 minutes
Close book and closed notes,
except a SINGLE piece of paper as a cheat sheet.

NO use of any device except calculators.

- This exam has 7 pages, including this cover page. Do all your work on these exam sheets. NO EXTRA PIECES OF PAPER WILL BE ALLOWED.
- Cross out all the scratch work that you do not want to be counted as part of your answer before you submit the exam.
- Be *specific, clear,* ***concise*** in your answers, and *explain* your answers.
- When the answer to a problem is not immediately clear, do not simply dump everything, relevant or irrelevant, on the paper. Irrelevant answers may lead to point-deduction as they show the lack of understanding of the problem.

Your name: _____

Student ID: _____

Problem 1 (20 points)

1.1 Circle zero or several application-layer protocols that use only TCP as their transport layer protocols?

(a) HTTP 1.1/2 (c) SMTP (e) BitTorrent (g) MPEG/DASH
(b) QUIC (d) IMAP/POP3 (f) DNS (h) Skype/VoIP

1.2 Circle zero or several application-layer protocols that are stateful?

(a) HTTP 1.1/2 (c) SMTP (e) BitTorrent (g) MPEG/DASH
(b) QUIC (d) IMAP/POP3 (f) DNS (h) Skype/VoIP

1.3 Circle zero or several statements that are **TRUE** for a peer-to-peer system?

(a) All systems always need to be on
(b) Transferring a file is faster than an equivalent client-server architecture
(c) They are not as scalable as client server architecture
(d) Are easier to implement than client-server systems

1.4 Circle zero or several pieces of information one **CANNOT** get by looking at an HTTP request message?

(a) Name of the web-page (c) Server's port number (e) Requester's IP address
(b) Server's host name (d) Server's IP address (f) Full URL of the request

1.5 Fill in the blanks:

The network protocols (and protocols in general) define

rules,

communication, and responses.

The most common HTTP method types are

GET and Host.

DNS protocol is a highly available database because DNS zone information (resource records) can be

cached and fetched iteratively.

HTTP protocol can scale because WEB data can be

Cached and pushed.

The common function (at least one) between TCP and UDP transport-layer protocol is

multiplexing.

In addition to this function, TCP also provides

reliable data transfer, congestion control and flow control & in order delivery

Problem 2 (20 points) Two hosts A and B are connected by a link with bandwidth of 1 Mbps (10^6 bits-per-second) and propagation delay between A and B is 1 millisecond. Host A has a 500,000-bit file to send to host B. A uses GoBackN reliable transport protocol and divides the file into 10,000-bit packets. The GoBackN protocol uses a fixed window size of 4 packets. You may assume the *transmission time* of ACK packets is negligible and no data or ACK packet ever gets lost.

Go back N = resend window if NAK/fail.

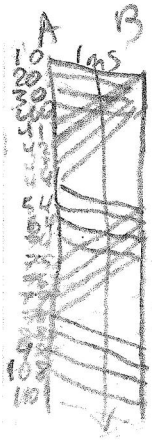

2.1 (6 points) How long will it take before the 12th packet has completely arrived at Host B? (drawing a diagram may help answer this question).

for ACK

Transmission delay: $\frac{L}{R} = \frac{10000 \text{ bits}}{10^6 \text{ bits/sec}} = 10$ ms

Prop delay = 0.0015 · 2 = 0.0025 = 2 ms

10ms + 2ms = 12ms total for 1 packet

~~To send 4 packets:~~

$\frac{10 \text{ ms}}{\text{pack}} \cdot 4 \text{ packets} = 40$ ms

110 ms

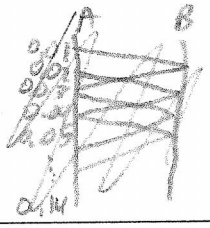

2.2 (6 points) How long will it take before the entire file is received by Host B?

$500{,}000 \text{ bits} \cdot \frac{1 \text{ packet}}{10000 \text{ bits}} = 50 \text{ packets} \cdot \frac{1 \text{ window}}{4 \text{ packets}} \approx 12.5 \text{ windows}$

From previous problem, takes 44 ms to send 1 window. Takes 22 sec for last window to send. So $\frac{44 \text{ ms}}{\text{window}} \cdot 12 \text{ windows} + 22 \text{ ms}$ (full)

$\frac{44 \text{ ms}}{\text{window}}$

= 550 ms

2.3 (6 points) How long will it take before the entire file is received by Host B if propagation delay is increased to 100 milliseconds?

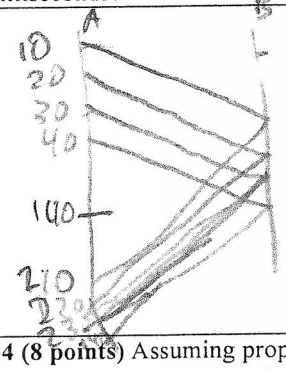

10 ms to load, 100 ms to send propagate. = 110 ms

240 ms to send 1 window & back

140 ms to get at B

$\frac{240 \text{ ms}}{\text{window}} \cdot 12 \text{ windows} + 140 = 3020$ ms

2.4 (8 points) Assuming propagation delay stays 100 milliseconds, is there a way for the file to be delivered to the host B faster by adjusting the window size? If so, what is the minimal window size that would allow the file to be received at B with shortest possible time (assume no other settings are changed)?

To get the shortest time, we need window to be big enough so that it captures last packet. To do this we need window size of 10 packets.

$\frac{0.01 + 0.01(3 \cdot 2)}{0.01}$

Problem 3 (20 points) Consider the following environment with a local DNS caching resolver and a set of authoritative DNS name servers.

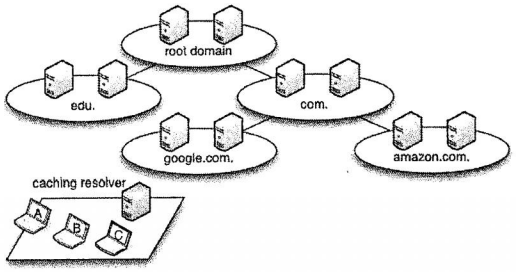

Assume that initially,

- the caching resolver cache is **empty**,
- TTL values for all records is **1 hour**,
- RTT between stub resolvers (hosts A, B, and C) and the caching resolver is **20 ms**,
- RTT between the caching resolver and any of the authoritative name servers is **150 ms**
- There are **no packet losses**
- All processing delays are **0 ms**

3.1 (4 points) At T=0 min, Host-A sends a query for "A record for amazon.com", and after receiving the answer sends a query for "A record for www.amazon.com". How long did it take to receive all the answers?

1st query: $T_{cache} + T_{root} + T_{com} + T_{amazon} = 20 + 150 + 150 + 150 = 470$ ms

2nd query: $T_{cache} + T_{amazon(www)} = 20 + 150 = 170$ ms

$470 + 170 = 640$ ms

3.2 (3 points) At T=40 min, Host-B sends a query for "MX record for google.com" that returns

```
google.com.          3600   IN    MX    10 primary.google.com.
google.com.          3600   IN    MX    30 backup.google.com.
primary.google.com.  3600   IN    A     74.125.28.27
backup.google.com.   3600   IN    A     173.194.211.27
```

(Similar to NS records, the DNS server may return "glue" A/AAAA records in addition to the requested MX records.)
How long did it take to get the answer? still in cache for 3.1 queries

$T_{cache} + T_{com} + T_{google} = 20 + 150 \cdot 2 = 320$ ms

3.3 (5 points) At T=70 min, Host-C sends a query for "AAAA (IPv6) record for mail.google.com", following at T=75 mins with a query for "AAAA (IPv6) record for hangout.google.com". How long did it take for Host-C to receive each of the answers (i.e., relative to T=70min for the first, and relative to T=75 mins for the second)? past TTL of 3.1, 3.2 & 3.4 for google.com

1st: $T_{cache} + T_{mail.google} = 20 + 150 = 170$ ms

2nd: $T_{cache} + T_{hangout.google} = 20 + 150 = 170$ ms

3.4 (5 points) List DNS records that the caching resolver has at T=90 minutes

google.com
hangout.google.com
mail.google.com
~~(+ those in 3.2 as well?)~~
primary.google.com
backup.google.com

3.5 (3 points) At T=100 minutes, all the authoritative servers for **.com** go offline. Circle the domain names that can be resolved by Host-A?

~~(a)~~ www.google.com (b) hangout.google.com (c) doc.google.com
~~(d)~~ www.amazon.com ~~(e)~~ video.amazon.com ~~(f)~~ aws.amazon.com

* Let NS = not set

Problem 4 (20 points) The following diagram shows a sequence of TCP packets for a client/server from your project 1, which include some of the sequence, acknowledgement numbers, and flags.

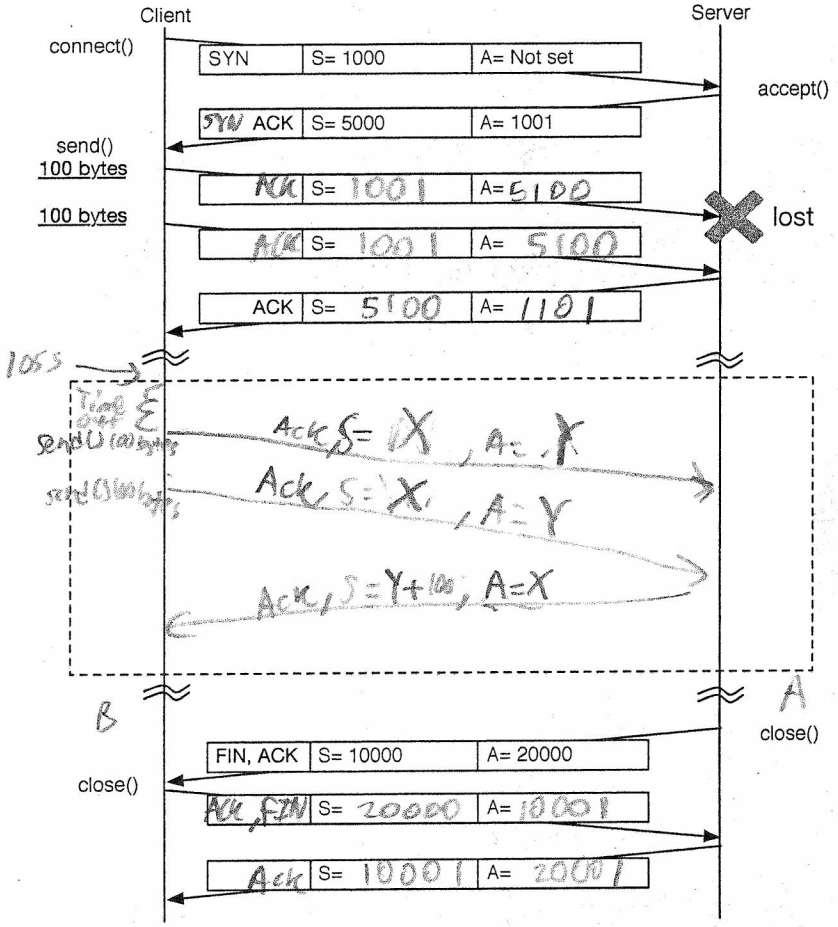

4.1 (10 points) In the figure above, fill in all the missing values for sequence, acknowledgement numbers and flags (SYN, ACK, FIN). For acknowledgement number write "Not set" if acknowledgment flag not set.

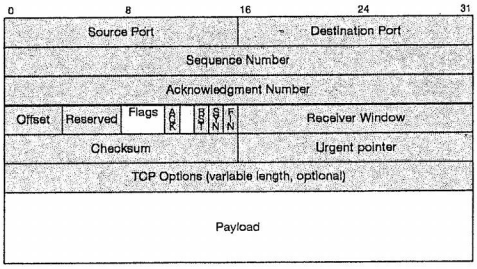

4.2 (5 points) One of the packets got lost. In the dotted box above, add the missing exchanges between the client and the server just after the loss has been detected. In the exchange, include flags, sequence number, and acknowledgement number (if applicable).

4.3 (5 points) What is the theoretical maximum of the TCP pipeline?
For a link with 500ms round trip delay imaginary 1000 Tbits/s link bandwidth, what is the maximum throughput that TCP protocol can sustain for that link (assume maximum packet size is 1000 bytes, hosts have infinite amount of buffering memory)?

$T = 10^{12}$

Theoretical max: 2048 bytes

$1000 \text{ bytes} \cdot \frac{8 \text{ bits}}{\text{byte}} \cdot \frac{s}{1000 \cdot 10^{12} \text{ bits}} = 8 \times 10^{-9} s$

$\frac{500 \times 10^{-3} s}{8 \times 10^{-9} s} = 6.25 \times 10^{7}$

Problem 5 (20 points): The following diagram shows the UDP packet header and HEX value of one of the captured UDP packets.

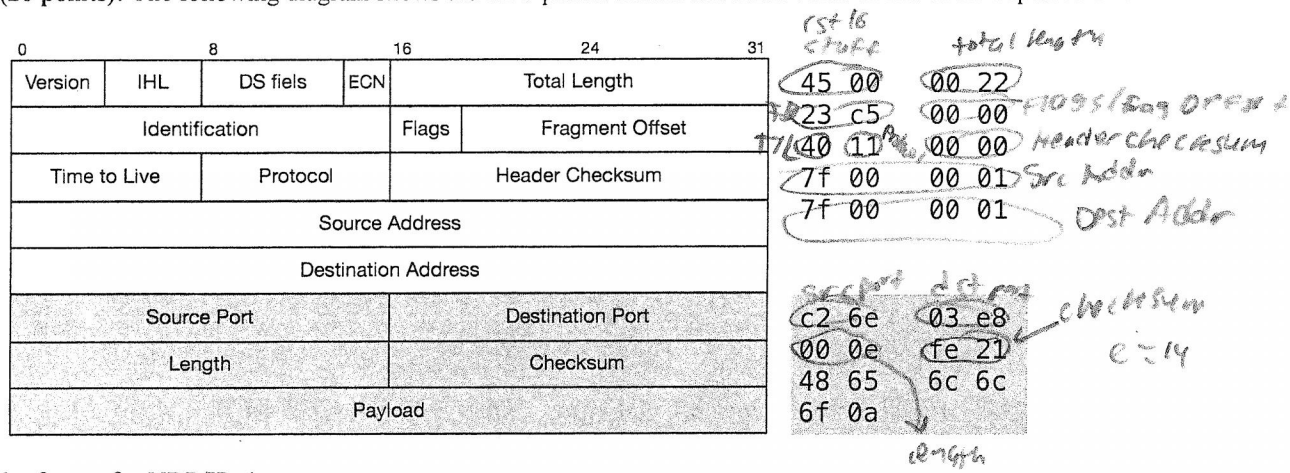

Pseudo-header format for UDP/IPv4

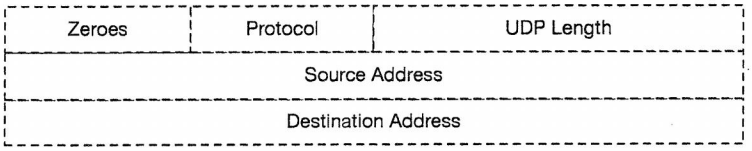

5.1 (5 points) Check correctness of UDP checksum. If it is incorrect, what should be the correct checksum?

checksum = fe 21 ⇒ $16^3 \cdot 15 + 16^2 \cdot 14 + 16 \cdot 2 + 1 = 65057$

Incorrect. Header checksum = 0, should be above

5.2 (5 points) Check correctness of IPv4 checksum. If it is incorrect, what should be the correct checksum?

length = 0e = 14 45 + 65 + 6c + 6c + 6f + 0a =

checksum = fe 21

5.3 (6 points) Please describe how this packet can be delivered to the destination application (i.e., how OS de-multiplex this packet) and on which port number this application should be listening on.

At the sender, we multiplex & handle data from the sockets & add transport headers for demultiplexing. At receiver, use header info to deliver received segments to correct socket. Host receives IP datagrams each w/ 1 transport layer segment, use IP addr & Port # to direct segment to appropriate socket

5.4 (4 points) Finish the following statements about UDP protocol:

The maximum size of a UDP payload is _____

The range for UDP port numbers is 0 - 55600

For a computer with two IP addresses (e.g., one for wireless and one for wired), there could be _____ maximum number of distinct UDP server applications. 2

To prevent anybody else to start a UDP server application, one need to start at least _____ number of applications, each creating one socket, binding, and listening on a single port.

Problem 6 (20 points)

6.1 (4 points) Assume that you want to send a secret message over email using PGP/GPG to a person you just googled on the Internet (you found his email and have a secret question to ask). Will you be able to do that? If yes, how (conceptually), if no, why?

Yes it is possible. One only needs to find the person's public key & encrypt the message to send with it, and attatch a digital signature (cerificate) as well. The person can decrypt with his public key

6.2 (4 points) Let's say you sent an email to the professor. If you haven't used PGP/GPG, he will not be able to know for sure that it came from you. List **at least two** facts that he can learnt from the received email that the sender couldn't fake. What can you do to ensure that the email is from you, including any **out-of-band** process that may be needed.

1. Use certificate to certify authenticity you are the indeed sender
2. Check the email address to see if email matches

6.3 (4 points) Let's say you go to a website over HTTPS protocol and get a warning that something is wrong with the certificate and browser rejects to proceed. List **at least four** reasons what can be wrong with a brief explanation what could have happened.

The certificate is invalid so the line is still ~~secured~~ encrypted but it could possibly not really be the website you're attempting to visit, say google.com. If you were to continue, you could be at a malicious site using a copy of google.com but different IP addr & if you input your info/passwords, it will send to the owner of the fake site, so in reality, you're sending data that's encrypted to a malicious person.

6.4 (4 points) Your professor travels a lot and whenever he has a chance he issues a DNS query for "A record for youtube.com". So far, he collected quite a bit of a collection of different responses. List **at least four** reasons why he gets different responses.

1. He's in a different geographic location, so will get query from closer DNS servers to him
2. Congestion/traffic in different cities might redirect him else where
3. ISP costs
4. Load balancing of DNS nodes

6.5 (4 points) HTTP/2 supports multiple streams and proactive push of data by the server. Give at least 2 reasons why people decided to develop QUIC.

1. To remove head-of-line blocking problems from earlier
2. Congestion control
3. priority streams
4. Runs over UDP