Network Systems (201600146/201600197), Test 1

February 17, 2017, 13:45–15:15

- This is an open-book exam: you are allowed to use the book by Peterson & Davie and the reader that belongs to this module, and the handout about peer-to-peer communication (i.e., the part of the Kurose&Ross book distributed via Blackboard). Furthermore, use of a dictionary is allowed. Use of a simple (non-graphical) calculator is allowed.
- Other written materials, and laptops, tablets, graphical calculators, mobile phones, etc., are not allowed. ***Please remove any such material and equipment from your desk, now!***
- Visiting the toilet without explicit permission of the supervisor is not allowed. During the last 30 minutes of the exam, no toilet visits are allowed.
- Write your answers on this paper, in the provided boxes, and hand this in.
- Total number of pages: 6.

---

1. Protocols and performance

Consider a link with a bandwidth of 1 Mbps, spanning a distance of 1200 km, with signals propagating at 200 000 km/s, and packets of 1000 bits.

2 pt (a) Compute the propagation time of a packet.

2 pt (b) Compute the transmission time of a packet.

2 pt (c) Suppose a sliding window protocol is used; what is the minimum Sending Window Size that will suffice to fully use the link? Explain.

2 pt (d) How many bits should the sequence number have, at least, for this SWS, with RWS=1, assuming packets cannot overtake each other? Explain.

1 pt (e) Why would one want to choose RWS > 1 ?

A. To prevent packets from overtaking each other.
B. To prevent problems if packets overtake each other.
C. To deal with variations in the bandwidth$\times$delay product.
D. In order to reduce the number of retransmissions if a single packet is lost.
E. In order to reduce the number of retransmissions if an entire window of packets is lost.

3 pt (f) Suppose we use a sliding-window protocol with RWS=1, SWS=2 and a sequence number space of 2, on a link where packets can be lost but cannot overtake each other. Sketch a time-sequence diagram in which a packet is delivered to the upper protocol layer twice.

Continued on next page...

2. Information theory and error-correcting codes

The speech of a newly elected president needs to be sent over a cable to reach the people who did not attend his inauguration. For simplicity, assume the president used only the following words, with their frequency of occurrence:

`great`	30%
`again`	30%
`absolutely`	19%
`fantastic`	19%
`peace`	2%

3 pt (a) Down to how few bits (on average) per word can his speech be compressed? Show your calculation.

1 pt (b) Suppose the order of the words in his speech is not completely random; either he *never* immediately repeats a word (i.e., after "great" he'll never immediately say "great" again), or he *always* says every word twice in a row. Is the resulting Shannon entropy per word higher or lower than in the completely random case?

A. Higher in both cases.
B. Lower in both cases.
C. Higher in case of no repetition, lower in case of doubling.
D. Lower in case of no repetition, higher in case of doubling.

2 pt (c) Which of the following codes is/are suitable for encoding the president's words into less than 2.25 bits per word on average?

word	probability	code A	code B	code C	code D	code E
`great`	30%	0	000	00	000	11
`again`	30%	1	01	01	001	10
`absolutely`	19%	00	10	10	010	01
`fantastic`	19%	01	11	110	011	001
`peace`	2%	110	001	111	1	0000
Suitable?						

In each of the boxes, write "yes" or "no"; or leave open if you don't know. Each correct answer gets +0.4 points, each wrong one –0.4 points.

1 pt (d) In his speech, the president makes an error. Can such an error be corrected by using an error correcting code for transmitting the speech?

A. Yes, but only if a Turbo code is used.
B. Yes, an error correcting code will certainly correct the error.
C. Yes, but the error is only corrected with probability arbitrarily close to 1, not certainty.
D. No, an error correcting code cannot correct errors that were already present at the source.
E. No: presidents don't make errors, they just present alternative facts.

Consider a binary symmetric channel with a raw speed of 100 bits/s and an error probability of 5%.

2 pt (e) Compute this channel's Shannon capacity.

1 pt (f) Suppose we want to send 70 bits/s over the above channel, and we can tolerate 0.01% bit errors.

A. This can be done even without using an error correcting code.
B. This needs an error correcting code, but it doesn't matter which code is chosen.
C. Only a suitably chosen error correcting code can do this.
D. This is not possible, no matter what error correcting code is used.

Continued on next page...

3. Applications

Consider a web server and a web client. The client (web browser) first downloads an HTML file of 1000 bytes; the HTML file then tells the browser to fetch an image of 10 000 and a movie of 10^7 bytes. The link speed is 1000 kbit/s, and the RTT is 2 ms. Assume the underlying protocols do not limit the transfer speed, i.e., their windows are large enough.

3 pt (a) How long does it take to load the page (including the image and the movie) using HTTP 1.1 with persistent connections and pipelining? Explain.

1 pt (b) How much longer would this have taken using HTTP 1.0 without parallel connections?

A. Equally long.
B. 1 RTT longer.
C. 2 RTTs longer.
D. 3 RTTs longer.
E. (Approximately) twice as long.
F. (Approximately) three times as long.

1 pt (c) Should the browser postpone playing the movie until it has been downloaded entirely?

A. Yes, movie playing software can't deal with an incomplete file.
B. Yes, otherwise only the top part of the picture can be shown.
C. No, but the earlier it starts playing, the more risk of stalling(*).
D. No, but the later it starts playing, the more risk of stalling(*).
E. No, delaying more than 400 ms would be disruptive for conversations.

(*) stalling means that the playback is temporarily stopped while waiting for data to arrive

Next assume the same movie is distributed using a peer-to-peer system. All computers involved have the same upload and download speeds as above (i.e., 10^6 bit/s). Initially, one computer has the file, and there are 100 computers ("peers") which want to get a copy of it.

3 pt (d) Calculate how long it takes for all peers to download the file; you may ignore the RTT for now.

1 pt (e) Which is/are the limiting factor(s) for the speed in this case?

A. None.
B. The server's upload speed.
C. The peers' download speed.
D. The peers' upload speed.

1 pt (f) Consider the IMAP protocol for e-mail. The first step in this protocol is an authentication step. Does this mean that if we make the use of IMAP obligatory, it is no longer possible to send messages with fake sender addresses (as demonstrated in the lecture) ?

A. Yes, since the sender's MTA could check that the sender address matches the IMAP authentication.
B. Yes, since the receiver's MTA could check that the sender address matches the IMAP authentication.
C. No, since sending mails is not done using IMAP but using SMTP.
D. No, since this authentication step happens after the message has been delivered to the receiver's MTA.
E. No, but building a wall between the USA and Mexico would solve the problem.

End of this exam.