Network Systems (201300179/201400431), Test 1

February 13, 2014, 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!***
- Although the questions are stated in English, you may answer in English or Dutch, whichever you are more comfortable with.
- You should always explain or motivate your answers, with so much detail that the grader can judge whether you understand the material; so just saying "yes" or giving a formula without explanation is not enough.
- 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.

---

1. A transatlantic link and reliable data transfer

A new cable has just been installed on the bottom of the Atlantic ocean. It can transport data at 20 Gbps. The cable is 10 000 km long, and signals propagate in it at a speed of $2 \cdot 10^8$ m/s. Data is sent in the form of packets that are each 40000 bits large.

4 pt (a) For a single packet, calculate the transmission delay and the (one-way) propagation delay.

4 pt (b) We want to send a file of $8 \cdot 10^9$ bits over this link using the stop-and-wait with 1-bit sequence numbers (also called "alternating bit") protocol. How long does this take?

3 pt (c) In order to speed things up, we switch to a sliding-window protocol with a larger send window. How large should that send window at least be for fully utilizing the link?

3 pt (d) How long does the file transfer take with that large window (calculated in the previous question)?

3 pt (e) Does it make sense to make the receive window larger than the send window? Why?

4 pt (f) Now assume that we use this transatlantic link to load a web page consisting of a small HTML file and one small image. How much time is save by moving from HTTP 1.0 (without pipelining, parallel or persistent connections) to HTTP 1.1 (with all of those)? Explain.

4 pt (g) Finally, consider a sliding-window protocol with Send Window Size = 1, Receive Window Size = 2, and 1-bit sequence numbers. Sketch a time-sequence diagram showing how things can go wrong if a packet gets dropped.

Continued on next page...

2. Information theory and error-correcting codes

A well-known organization with a three-letter name is spying on e-mail. For each e-mail, it observes whether the e-mail is *harmless* (80 % of the cases), *suspicious* (19 %), *terroristic* (0.9 %), or *from Syria* (0.1 %)

4 pt (a) How much information is this, per observation?

4 pt (b) Propose a way to encode the observations in less than 1.5 bits on average, and show that your code indeed achieves this.

The organization sends its observations over a binary symmetric channel, with an error probability of 0.1% and a raw speed of 1000 bits per second.

4 pt (c) How many observations can be sent per second, assuming very good coding, if the error rate should be made extremely low?

4 pt (d) A CRC is normally used to *detect* errors. Could a CRC also be used to *correct* a bit error? How, or why not?
Hint: in your answer, take into account the fact that CRCs are guaranteed to detect every error of 1 or 2 bits.

---

3. Peer-to-peer applications

Consider a peer-to-peer system as follows. There is one server, denoted S_F, that makes available a large file $\mathcal{F}$ of size F bytes to n peers. The upload rate from server S_F is u_F bytes per second. The download rate of each of the n clients is d bytes per second; the upload rate of each peer is u bytes per second.

For this scenario, it is known that the overall download time D_F for the file $\mathcal{F}$ is lower bounded as follows:

$$D_{\mathcal{F}} \geq \max\left\{\frac{F}{u_F}, \frac{F}{d}, \frac{F}{\frac{u_F}{n} + u}\right\}.$$

5 pt (a) Explain this expression, including the $\geq$-sign, the max-operator, and the three subexpressions within the curly braces.

4 pt (b) Now, consider the scenario that there is also another server active, denoted S_G, that makes available another large file $\mathcal{G}$, of size G, again to the same n peers. Give an expression for the minimum download time for both files $\mathcal{F}$ and $\mathcal{G}$, in a similar style as given above (for just $\mathcal{F}$), that is:

$$D_{\mathcal{F} \text{ and } \mathcal{G}} \geq \max\{\ldots, \ldots, \ldots, \ldots\}.$$

3 pt (c) In distributed hash-tables (like used in Chord), a distributed database is maintained for storing the overall set of (key, value)-pairs on all available machines. Storing such a table in just one place (on one server), would make things much simpler. Give two reasons why this is not done?

End of this exam.