Computer Networks
25th of June 2009

- This exam consists of 8 questions with subquestions. Every subquestion counts for 10 points (except the last one which counts for 1 point of your final grade).
- Mark every page with name and student number.
- Use of books, additional course material or electronic devices is prohibited.
- Always explain your answers. At the same time, keep your answers short and to the point. Do not use pencil or red ink.

---

1. **Multiple choice**

   Please answer the following questions. For each question, mark the correct answer. There is exactly one correct answer per question. The full question is worth 10 points: you are awarded 2 points for each correctly answered subquestion. However, each wrongly answered question results in 2 negative points. Don't answer at random since this will decrease your score. Concentrate on the questions you can answer with certainty.

   i. Consider a link-layer protocol that implements byte-stuffing using 'e' as the escape character, and 'f' as the frame delimiter. How would the message frame:

   four apples and five pears

   be transmitted on the link?

   (a) efour applees and efivee peears
   (b) ffour apples and efive pearsf
   (c) effour applfes and ffivfe pfears
   (d) fefour applees and efivee peearsf
   (e) None of the above

   ii. Which of the following protocol layers is not explicitly part of the IP reference model?

   (a) Application
   (b) Physical
   (c) Transport
   (d) Session
   (e) Network

   iii. Which of the following statements about circuit-switched networks is true?

   (a) Circuit-switched networks make more efficient use of network resources than packet-switched networks
   (b) Circuit-switched networks do not require a connection-setup step for communication.
   (c) Bandwidth in a circuit-switched network can be shared between users via time-division multiplexing
   (d) Circuits can't provide strong performance guarantees
   (e) The Internet is a circuit-switched network at the Network Layer.

   iv. Which of the following protocols is stateless?

   (a) SMTP
   (b) Telnet
   (c) HTTP 1.0

v. Which of the following statements about DNS is true?

(a) The TTL value in DNS records makes it possible to change the mapping of a hostname to an IP address (e.g., for load balancing). However, when we map the name to a new IP address, all existing connections to the old IP address have to be moved too.

(b) It is important that information we receive from DNS is reliable. We therefore use TCP for querying the name servers.

(c) Due to their iterative or recursive lookups, the first DNS query for a particular name is fairly slow. However, caching can be applied at all levels to make subsequent queries for the same name faster.

2. **Physical**

(a) Briefly explain the difference between frequency division multiplexing (FDM) and time division multiplexing (TDM).

(b) Code division multiple access (CDMA) is used between 3 stations with the following chip sequences (shown both in binary and bipolar form):

```
A: 0 0 0 1 1 0 1 1          A: (-1 -1 -1 +1 +1 -1 +1 +1)
B: 0 0 1 0 1 1 1 0          B: (-1 -1 +1 -1 +1 +1 +1 -1)
C: 0 1 0 1 1 1 0 0          C: (-1 +1 -1 +1 +1 +1 -1 -1)
```

We receive the `(0 0 -2 +2 0 -2 0 +2)`. →Who sent what bits?

(c) We are trying to transmit a single byte with a bit rate of 2400 bps.

i. What is the period T of the first harmonic?

ii. If your link supports a maximum frequency of 3000 Hz, how many harmonics can we have?

(d) What is the purpose of modulation?

3. **Datalink: multiple access protocols.**

We studied a number of multiple access protocols in this course, including (1) TDM (time division multiplexing), (2) CSMA (carrier sense multiple access), (3) Slotted Aloha, and (4) Token passing. Suppose there are **N** stations on a LAN that has capacity (transmission rate) **C**. All packets have a fixed length **L** and the end-to-end propagation delay of the channel is **P**. For each of the protocols above, answer the following questions:

(a) Suppose only one station ever has a message to send (i.e., the other N-1 stations generate no traffic). →What is the maximum possible throughput seen by this single node under each of the protocols above?

(b) Suppose now that all stations have the same average traffic arrival rate. We are interested in the aggregate throughput of the LAN. →For each of the above protocols, is it possible to achieve a throughput of 1 (i.e., have the channel always be fully utilized)? If not, indicate how/why the protocol limits the maximum throughput to less than 1.

(c) In a heavily loaded network, what is the worst case amount of time a node has to wait under each of the protocols, before it can send a message?

4. **NAT**

(a) NAT helps us deal with shortage of IP addresses by mapping multiple local IP addresses onto a single IP address. Also, TCP port numbers are mapped. "Great", someone says, "All client machines in a domain can simply be grouped together and use local addresses in the 10.0.0.0/8 address range, so that we can have 2^{24} clients accessing the Internet (e.g., google, amazon, cnn, etc.) *at the same time* while using only a single IP address." Assume these machines are indeed clients only. → Explain what is wrong with the argument.

5. **TCP**

A man in a mid-life crisis goes sailing across the Atlantic. To stay in touch with his students he purchases a shortwave radio which he uses to carry Internet connections. All is well, except that his TCP connections suffer from the large number of lost packets. Sometimes the packet loss occurs as a burst of lost packets.

(a) Explain what will happen to the TCP connections if there is frequent packet loss by discussing TCP congestion control.

(b) Can you think of a solution to improve the connections without modifying TCP?

6. **Performance**

   Define network bandwidth and network latency. For many applications (file transfer, web browsing, BitTorrent, etc.) increasing the available network bandwidth will increase performance. However, for some communication-intensive applications providing more bandwidth will not improve performance, even though both endpoints of the communication spend most of the time waiting for messages to arrive on the network. Explain why.

7. **TCP**

   (a) In 'standard' TCP, the congestion control mechanism drops the window size to 1 MSS (maximum segment size) whenever an acknowledgment does not arrive within the time-out period. Why?

   (b) In some cases, such a reduction of the window size is rather awful for the performance. Give an example where this is the case.

8. **The question question**

   Give an interesting question of your own about the course material (*and* include the answer). 'Knowledge' questions (questions that aim at reproducing some material from the course material directly) contribute 0.5 points to your final grade, while 'insight' questions contribute a maximum of 1.0 points to your final grade. In both cases, the answers have to be correct,