Computer Networks

28th of January 2011

- This exam consists of 6 questions with subquestions. Every subquestion counts for 10 points.
- Mark every page with name and student number.
- Use of books, or additional course material is prohibited. You may use a calculator..
- Always explain your answers. At the same time, keep your answers short and to the point. Do not use pencil or red ink.
- Answers can be in English or Dutch.

---

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 1 point for each correctly answered subquestion. However, each wrongly answered question results in 1 negative point. Don't answer at random since this will decrease your score. Concentrate on the questions you can answer with certainty. *If needed, you are allowed to add one line of explanation to each answer - but no more.*

   i. Two hosts simultaneously send data through a link of capacity 1Mbps. Host A generates data with a rate of 1Mbps and uses TCP. Host B uses UDP and transmits a 100bytes packet every 1ms. Which host will obtain higher throughput?

      (a) Host A.
      (b) Host B.
      (c) They obtain roughly the same throughput.
      (d) They experience congestion collapse and negligible throughput.

   ii. What is the theoretical upper-bound on the number of simultaneous TCP/IPv4 connections that a host with a single IP address can handle? (Ignore memory limitations.)

      (a) 2^{16}
      (b) 2^{32}
      (c) 2^{64}
      (d) unlimited

   iii. Host A sends a TCP segment (Seq = 43, ACK = 103), to which host B replies with a TCP segment (Seq = 103, ACK = 57). The payload of the first TCP segment is

      (a) 14 bytes long.
      (b) 43 bytes long.
      (c) 46 bytes long.
      (d) 57 bytes long.
      (e) 60 bytes long.

   iv. Longest prefix matching is used:

      (a) in routers to know on which link interface to forward packets.
      (b) in classless addressing to use the address space more efficiently than in classful addressing.
      (c) by NAT to increase the available address space in home networks.
      (d) to assign subnet masks.

(e) None of the above.

v. An authoritative DNS server knows a top-level domain server via:

(a) hostname.
(b) IP address.
(c) canonical hostname.
(d) alias.
(e) domain.

vi. A user requests a web page that consists of some text and 3 images. The browser's cache is empty. For this page, the client's browser:

(a) sends 1 http request message and receives 1 http response messages.
(b) sends 1 http request message and receives 3 http response messages.
(c) sends 1 http request message and receives 4 http response messages.
(d) sends 3 http request messages and receives 3 http response messages.
(e) sends 4 http request messages and receives 4 http response messages.

vii. Queueing on a router's input port occurs mainly when:

(a) all packets arriving on the input port are destined for the same output port
(b) transmission rate of the output port is much higher than that of the input port
(c) all 'head of line' packets (the packets in the front of the queue) in all input queues are destined for different ports
(d) multiple 'head of line' packets in multiple queues are destined for the same output port and the queues are serviced in a FIFO fashion
(e) multiple 'head of line' packets in multiple queues are destined for the same output port and the queues are serviced in a LIFO fashion

viii. To benefit from IPv6

(a) all the routers in the Internet between source and destination should support IPv6
(b) we also need new routing algorithms
(c) we need a new type of ICMP to 'probe' the optimal MTU size
(d) we generate an icmp error message back to the sender whenever our datagram is too large for a network link (greater than the MTU)

ix. Defining efficiency in terms of resource utilisation, CDMA (Code Division Multiple Access) with a chipping sequence length of 16

(a) will be less efficient than FDMA with 16 frequency bands if the number of senders is less than 16
(b) will be as efficient as FDMA with 16 frequency bands but not as efficient as TDMA with 16 slots
(c) will be as efficient as FDMA and TDMA, but much less complicated to implement
(d) None of the above / you cannot tell

x. Manchester encoding is principally designed to:

(a) ensure that a transition occurs in the centre of each bit period.
(b) have more than one symbol per bit period.
(c) increase the bandwidth of a signal transmitted on the medium.
(d) ensure that the line remains unbalanced.

2. **Datalink**

(a) Give the two main reasons why the 802.11 protocol does *not* use CSMA/CD?

(b) We have seen that IEEE 802.3 differs from regular Ethernet in that it does not have a *type* field. Instead, it uses those two bytes in the frameheader to indicate the *length* of the data. Discuss what the implications are for demultiplexing (e.g., is it harder, simpler, the same?). Explain.

(c) The USB protocol uses a CRC with Generator 100101. Calculate the CRC of the (data) bit string 1100011110110001

(d) With no more than 2^{32} addresses, addresses in IPv4 are quite scarce. We do not have this problem in Ethernet, where we have a whopping 2^{48} address to play with. Since collisions are no longer a problem in modern-day Ethernet and switches automatically learn the 'routes' in the network, would it be possible (in theory) to build a worldwide network (such as the Internet) with millions of users using just Ethernet (i.e., dropping IP altogether)? Why/why not? (If your answer is 'no', give the 3 main technical reasons why not. If your answer is yes, argue why using at least 3 arguments.)

(e) In the CSMA/CD protocol, which of the following conditions on the transmission delay T_{trans} and the propagation delay T_{prop} has to be satisfied to guarantee that a node always detects a collision?

(a) $T_{trans} > T_{prop}$
(b) $T_{trans} > 2T_{prop}$
(c) $T_{trans} < T_{prop}$
(d) $2T_{trans} > T_{prop}$

Explain why!

3. **Network/Routing**

(a) What are the four causes, or components, of delay in a packet switched network? More interestingly perhaps, what can be done to reduce each of these components?

(b) For the network shown in Figure 1 (page 4), indicate (in the tables) how the Distance Vector routing algorithm builds the corresponding distance tables for each node. Indicate which tables correspond to which nodes and add all the necessary entries in the tables. Show all the steps until the algorithm stops.

(c) Consider the network below, with the initial distance vectors towards A filled in for nodes B, C, and D. Node A goes down. Show very briefly how (pure) distance vector routing may lead to a count-to-infinity situation.

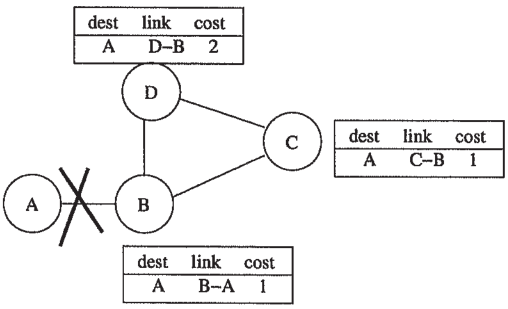

(d) Explain how *poisoned reverse* helps against the count-to-infinity problem.

(e) Using the network and initial distance vectors shown in the figure, describe a scenario where poisoned reverse does *not* solve the count-to-infinity problem. [Hint: you may want to consider the case where one node has already heard about the 'poisoned route', while the other (which hasn't yet) starts advertising before it hears about it.]

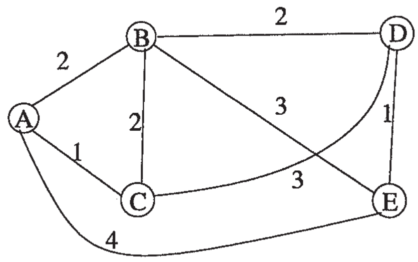

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Node	cost via				
	A	B	C	D	E
destination A					
B					
B					
D					
E					

Figure 1: Distance vector routing

4. **TCP.**

A lonely sailor on the Pacific Ocean wants to access his email from his server at home. The one-way latency is 100ms and the bandwidth is, on average, 1 Mbps.

(a) How large should the TCP receiver window be to achieve 1 Mbps (assuming ideal conditions)?

In the remainder of these questions, assume that after the TCP connection is established, TCP slow start will begin, that the initial congestion window is 1 segment, and that the initial threshold is 64 segments. All segments that carry data have the length of a maximum segment size, which we assume to be 10KB.

(b) Unfortunately, the wireless (radio) connection creates high error rates. One in every 8 segments is corrupted. Given the initial threshold value, will the TCP connection ever get out of slowstart? Explain.

(c) Assume that *exactly* the 8th outgoing segment (i.e., sent by the sailor) is lost. (i) At that point, what is the size of the congestion window (in MSS) if we assume that acknowledgments for segments 1-7 have all been received? (ii) What will be the new size of the threshold?

(d) How much *user data* (not including the IP and TCP headers) will have been transmitted by the sailor at that point?

5. **Wireless/mobile**

(a) Can we use the following two codes ("chipping sequences") for a CDMA protocol to avoid collision between two stations? Why or why not?

- c1 = `(1;-1;-1;-1; 1; 1;-1;-1)`
- c2 = `(1; 1;-1; 1; 1;-1; 1; 1)`

(b) What is triangle routing? Does mobile IP use triangle routing or direct routing?

(c) What is the main purpose of using RTS/CTS packets (Request to Send/Clear to Send packets) in wireless networks? Explain this with one example.

6. **The 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) may give you 5 points (max), while 'insight' questions may give you a maximum of 10 points (max). In both cases, the answers have to be correct.