

Computer Networks

19 December 2014

-
- 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, additional course material, and calculator 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. Give your answers on the exam paper (if needed, you may request additional paper.)
 - 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 2 point for each correctly answered subquestion. *If needed, you are allowed to add one line of explanation to each answer – but no more.*

i. Stop-and-wait

- (a) is a sliding windows protocol with a window size of 2: one for the first segment and one for the next
 - (b) does not really need any sequence numbers as it waits for the acknowledgment of every segment
 - (c) is bad for congestion as it does not allow the sender to reduce its rate
 - (d) is per definition not compatible with TCP-like slow start and additive increase
-
-

ii. Which of the following statements about routing and forwarding is true?

- (a) Routing algorithms like OSPF and RIP are not very well-suited for networks with rapidly changing topologies
 - (b) A connection-oriented communication service provides a virtual circuit across the network prior to communication in which resources (like bandwidth) are reserved on the intermediate hops
 - (c) As soon as a pure store-and-forward router receives enough bytes of a packet to determine the route, it starts transmitting it on the outgoing link. Thus, it may be that the last bytes of a packet have not yet arrived at the router, when the first bytes are already being forwarded
 - (d) In IPv4 routing, the router will look at the packet's destination address *and* its source address to determine the next hop
-
-

iii. Which of the following statements about protocols is true?

- (a) IP, the “narrow waist” of the Internet, is great because it allows higher layer protocols to use all the advanced features offered by any datalink layer.

- (b) It is possible to send an IP packet inside another IP packet (as long as we make sure that the protocol field of the outer datagram specifies that the payload is another IP datagram). In that case, the TTL value of the inner datagram will not be decremented at each hop.
 - (c) IPv6 has a variable-sized header length to allow for a variety of options
 - (d) The Ethernet checksum is different (and stronger) than that of IP, TCP, and UDP, but it is like the others in that it is added in a header—prefixed to the data.
-
-

iv. Which of the following statements about addresses is true?

- (a) ARP allows a host to find an IP address given a MAC address
 - (b) Suppose that when connecting to a web server, a NAT is used. This means that in the outgoing direction (towards the server), the IP source address, the TCP source port, and the TCP destination port will all be remapped.
 - (c) We can fit exactly four $\backslash 18$ networks in a single $\backslash 16$ network.
 - (d) After obtaining an IP address via DHCP, a host discovers its local DNS server by means of a broadcast.
-
-

v. Which of the following statements about application layer protocols is true?

- (a) Cookies in HTTP are stored at the web server, so they can be sent to the client regardless of the computer he/she uses to connect to the server
 - (b) SMTP, the simple mail transfer protocol, can be used for both sending and receiving of email, but we also use protocols like POP3 and IMAP for receiving.
 - (c) A recursive DNS lookup *always* contacts the root servers first
 - (d) If we register our DNS records with shorter TTL values, we will typically increase the load on our authoritative server as there is less benefit from caching.
-
-

2. Physical Layer

- (a) What is the formula for Shannon's channel capacity?
 - (b) The signal to noise ratio (SNR) is 30 dB. (i) How much stronger is the signal than the noise?, (ii) Assume that the channel has a bandwidth of 1000 Hz. Is a capacity (bit rate) of 10000 bits per second possible on this channel?
-
-
-
-

- (c) "CDMA is potentially more efficient than TDM and FDM, because all senders can transmit at

the same time and each can transmit at full capacity.” Explain whether you agree.

3. Datalink

- (a) Give two main reasons why the 802.11 protocol does *not* use collision detection.

- (b) On a USB connection we receive the following transmission (including the CRC bits): 110001111010000101110. USB uses a CRC with generator 100101. Discuss whether or not you think there may have been an error in the transmission.

4. Network layer and routing

- (a) Consider the following shortest path algorithm as used for finding a shortest path route from

S (source node) to D (destination node). Define N as the set of nodes in the connected graph, define M as the set of nodes selected as follows:

- 1) Start at a source node S, initialize set M to empty
- 2) Find a neighbor node not in set M with lowest link cost, add this node to set M
- 3) Repeat step (2) until D is in set M

Will this algorithm correctly find least-cost routes? If this is not a good algorithm, give an example where it fails.

- (b) In distance vector routing, we use poisoned reverse to prevent the count-to-infinity problem. (i) explain briefly how it works. (ii) Given an example of a situation where poisoned reverse does not prevent a count-to-infinity problem.

5. Sliding windows

- (a) Explain how Go-Back-N works.

- (b) In general, Go-Back-N solves the poor link utilisation of stop-and-wait. There are (pathological)

cases where this is not true. Describe a case where Go-Back-N performs even worse than stop-and-wait.

- (c) We have seen that stop-and-wait is very inefficient on links with a large bandwidth-delay (BD) product. Let us examine to what extent this is true. Assume we have a link with a high bandwidth B and delay D and also assume that ACKs have negligible size. For each of the following “improvements” say whether they improve the transmission speed. If so, also explain what the new maximum speed can be, and how it can (theoretically) be achieved. Assume the link speed is C (bits per second).
- (i) rather than a single connection, the sender opens many connections to the receiver
 - (ii) the sender sends very large packets

6. TCP.

The year is 2023 and a Dutch broadcasting company has finally established a permanent base on Mars, where the colonists are part of a new reality TV show much like Big Brother and will be under 24/7 video surveillance¹ It is a remote place (it takes 20 minutes to send a signal from Earth to Mars and back) and for the video transmissions from all the cameras, the organisers have planned a 10Gbps TCP connection, to be set up by the colonists.

In the remainder of these questions, assume that (i) *after the TCP connection is established, TCP slow start will begin*, (ii) *the 3rd segment of the connection establishment contains data*, (iii) *the initial congestion window is 1 segment*, and (iv) *the initial threshold is 64 segments*.

¹Your humble lecturer wishes to point out that he is not making this up. There really exists a Dutch organisation with plans to have a “Big Brother on Mars” in 2023. He does not think this is the best idea to come out of this country.

- (a) Assuming that initially there is no connection yet, after how many RTTs will the connection leave slow start?

The competing TV stations are jealous and decide to jam the transmission periodically. Let us assume that *exactly* the 25th segment is zapped by these guys when it is exactly midway between Mars and Earth (note: this could be an earth-to-mars or a mars-to-earth segment).

- (b) How much time has passed when the segment is zapped (rounded to the nearest minute) from the transmission of the first bit by the sender to the segment getting zapped?

This image shows a single sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

Back on earth... Consider Figure 1, which shows the plot of the TCP congestion window size as a function of time for two TCP connections A and B. In this problem we will suppose that both TCP senders are sending large files. We also assume that the packet loss events are independent in connection A and B. TCP connection A uses fast recovery.

- (c) (i) What are the values of the Threshold parameter between the 1st and 6th transmission rounds




Figure 1: Evolution of TCP's congestion windows for two connections A and B

for connection A. (ii) Same for connection B, but between the 1st and 14th transmission rounds. (→If you cannot determine the threshold, explain why not.)

At the 6th transmission round, connection A experiences a segment loss which is detected by a triple duplicate ACK. After that, it does not experience any loss until round 12, at which point a time-out occurs. After that, neither of the connections experiences any loss.

- (d) Draw (on the figure) the CongWin values of both connections up to the 20th transmission round. (Use space below for additional explanation and comments, if any)
