

Always explain your answers concisely and be sure to be to-the-point.

Part I

This part covers the same material as the midterm exam.

1a Why does the 802.3 (Ethernet) description also specify what the maximum segment length is? 5pt

802.3 describes a protocol with collision detection. If host A starts transmitting a frame, it should be able to detect a collision before it finishes its frame transmission. This means that any bit sent by another station should make it to A before A finishes. Bit propagation takes time, depending on the length of a segment. To limit the propagation time, you need to specify a maximum segment length.

1b Gigabit Ethernet allows for carrier extension and frame bursting. What do these techniques establish and why are they necessary? 5pt

Both techniques allow a sender to extend a frame to a minimum of 512 bytes, enabling collision detection for cable segments longer than 25 meters.

1c For a switched Gigabit Ethernet connection, there is no maximum segment length specified. Why not? 5pt

There is simply no need for collision detection, provided we can assume that the line operates in full-duplex mode. Mentioning the latter is important: in half-duplex mode, you will have contention so that collision detection is necessary.

2a If τ is the propagation speed of a signal, and ρ the transmission rate, how many bits can a wire of length L contain? 5pt

Assuming τ is expressed in meters per second (m/s), ρ is expressed in bits per second (b/s), and L is expressed in meters, then the total bit capacity of the wire is $\rho/\tau \cdot L$.

2b Consider a token ring network without any artificial delays, operating at a transmission rate ρ . Each computer in the ring introduces a delay of δ seconds. With a propagation speed over the wire of τ , and a token length of R bits, what is the minimum ring length? 5pt

If you take into account that computers may be switched off, then we need to ensure that the complete token fits on the ring. τ/ρ expresses how many meters you need to store a single bit, so that the minimum ring length equals $\tau/\rho \cdot R$. If you take into account that a computer always introduces a delay, then $\delta \cdot \rho$ expresses how many bits can be "stored" at a single computer, leaving $R - \delta \cdot \rho$ to be handled by the rest of the network. Therefore, with N computers, the minimum length is simply $L = \max\{\tau/\rho \cdot (R - N \cdot \delta \cdot \rho), 0\}$

3a What is the necessary and sufficient condition for constructing a k -bit error detecting code? 5pt

The Hamming distance between any two transmitted strings should be at least $k + 1$.

3b Consider the following 2-dimensional parity code (see the example below). A string of $b_1 b_2 \dots b_n$ of n bits is split into k parts of l bits, i.e., $n = k \cdot l$. The string is then represented in a $k \times l$ matrix, where the i -th row contains bits $b_{l(i-1)+1} \dots b_{li}$, and the j -th column contains bits $b_j b_{j+l} \dots b_{j+(k-1)l}$. The i -th row is extended with a 1 if that row contains an odd number of bits, and with a 0 otherwise. Likewise, the j -th column is extended with 1 or 0, respectively. Show by example that this is a 1-bit error-correcting code. 10pt

| original string | encoding | transmitted string |
|-------------------|-----------|-------------------------------|
| | 1 0 1 0 | |
| | 1 1 0 0 | |
| 1 0 1 1 1 0 1 1 1 | 1 1 1 1 | 1 0 1 1 1 0 1 1 1 0 0 1 1 0 0 |
| | ----- | |
| | 1 0 0 | |

You can simply take the example, and show what happens when 1 bit is flipped. Be sure to distinguish the original bits from the bits that have been added for error correction.

- 4 Limited-contention protocols dynamically adapt to traffic intensity. What problem do they solve by such an adaptation? 5pt

Broadly speaking, there are two classes of protocols: contention protocols and collision-free protocols. The first class allows senders to transmit data, but retransmissions are necessary when collision occur. The second class avoids collisions altogether by means of, e.g., tokens. Contention protocols (which are essentially optimistic) operate best when network traffic is low, while the opposite holds for collision-free protocols (which are, in fact, pessimistic). Limited-content protocols operate as contention protocols when traffic is low, but behave as collision-free protocols when there's a lot of traffic. The problem that they thus solve, is that they are not overly optimistic (pessimistic) when traffic is high (low).

Part II

- 5a Explain the difference between integrated and differentiated services as provided by the Internet IP layer. 5pt

Integrated services are services for supporting multimedia traffic. Their implementation is based maintaining a flow between communicating parties, for which reason explicit flow management is required. In practice, this means that routers need to maintain state, allocate buffers, and so on. Differentiated services have no notion of a flow, but distinguish only classes. Each packet belongs to a specific class. Routers take decisions on prioritizing packet processing by considering only the class to which a packet belongs.

- 5b Explain how fair queuing works in routers and which problem it solves. 5pt

Normally, routers simply pick up incoming packets and forward them to outgoing interfaces. In principle, all packets are treated equally. This approach is beneficial for hosts that send large packets, as they effectively consume more available bandwidth in comparison to hosts sending small packets. With fair queuing, all incoming packets are conceptually split into unit-sized subpackets. A router will then forward these subpackets to outgoing interfaces, treating subpackets as equals. However, only after all the subpackets of a packet have been queued for an outgoing link, will it be possible to actually transfer that packet. In effect, all hosts are now given an equal share of the available bandwidth. See also Figure 5-36.

- 5c Multi-Protocol Label Switching (MPLS) is popular for multimedia streaming in wide-area networks. Why? 5pt

MPLS adds a label to every packet. This label is used by routers to look up the outgoing interface for an incoming packet, where it is assumed that the lookup table has been configured in advance. In effect, MPLS allows for the setup of a virtual circuit between two hosts, including the allocation of resources. This pre-allocation of a route makes it much easier to guarantee quality of service, which is exactly what is needed for multimedia streaming.

- 6a What does the 3-way handshake in TCP establish? 5pt

It establishes that a client and server agree on the sequence numbers that are to be used in the packets that the client sends to the server, and those to be used in the packets sent from the server to the client.

- 6b When a client sets up a TCP connection, it may request a buffer size at the server that is larger than its congestion window size. Does this make sense? 5pt

Yes, it does. The congestion window merely indicates how much data the client can send in a single burst. It says nothing about the amount of data that the receiver is willing to buffer, so that the client and server can operate asynchronously. Of course, if the receiver buffer size drops below the congestion window size, the client can only send as much data as the receiver is willing to accept.

6c Explain what is meant by the silly window syndrome.

5pt

This phenomenon occurs when there is a TCP connection between two hosts, and when the receiving application systematically reads only a single byte from the receiver buffer. If the receiving party notifies the sender every time the receive buffer has a one-byte slot open, we may see a huge number of packets being transferred from client to server, each packet containing only a single byte of data. Obviously, this is very inefficient considering that the TCP and IP header alone already consume at least 40 bytes.

7 Below is the output of querying a DNS server for MX records for the vu.nl domain. What does it tell us?

5pt

```
seuss % dig MX vu.nl

; <<>> DiG 9.2.5 <<>> MX vu.nl
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 59633
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 3, ADDITIONAL: 7

;; QUESTION SECTION:
vu.nl.                IN      MX

;; ANSWER SECTION:
vu.nl.                60      IN      MX      20 mail2.vu.nl.
vu.nl.                60      IN      MX      10 mail.vu.nl.

;; AUTHORITY SECTION:
vu.nl.                86400   IN      NS      ns1.surfnet.nl.
vu.nl.                86400   IN      NS      star.cs.vu.nl.
vu.nl.                86400   IN      NS      ns.vu.nl.

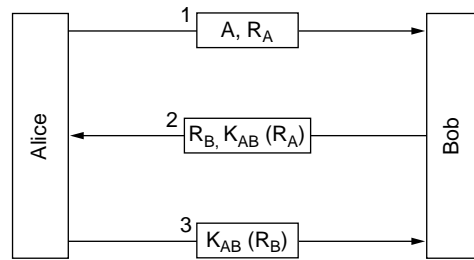
;; ADDITIONAL SECTION:
mail.vu.nl.           86400   IN      A       130.37.129.161
mail2.vu.nl.          86400   IN      A       130.37.129.165
ns.vu.nl.             86400   IN      A       130.37.129.4
ns1.surfnet.nl.       168     IN      A       192.87.106.101
ns1.surfnet.nl.       168     IN      AAAA    2001:610:1:800a:192:87:106:101
star.cs.vu.nl.        86400   IN      A       130.37.24.6
star.cs.vu.nl.        86400   IN      A       192.31.231.42

;; Query time: 322 msec
;; SERVER: 130.37.20.3#53(130.37.20.3)
;; WHEN: Fri May 27 11:43:48 2005
;; MSG SIZE rcvd: 255
```

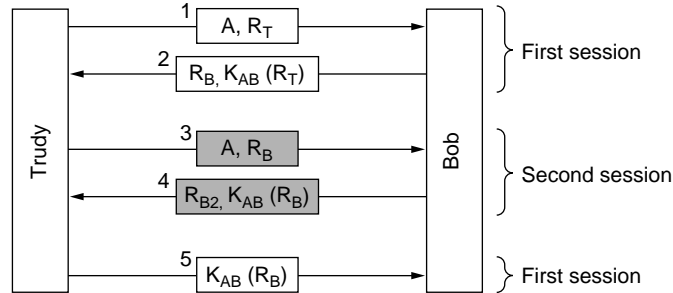
Lots of things. You should at least mention that the mail servers for vu.nl are mail2.vu.nl (lowest priority) and mail.vu.nl (highest priority). The IP addresses for these two mail servers are 130.37.129.165 and 130.37.129.161. Furthermore, there are three name servers responsible for this domain. Of these name servers, ns1.surfnet.nl has an IPv6 address (but only a single network interface). Server star.cs.vu.nl is a multihomed IPv4 host.

8a Show that the following protocol is subject to a reflection attack.

5pt

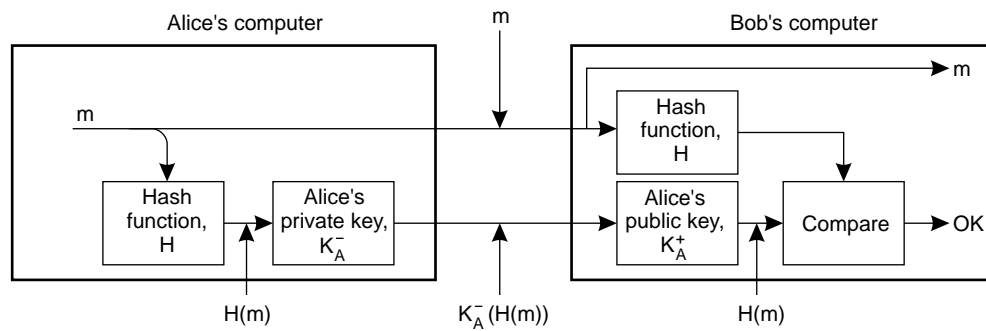


You should come up with the following picture, and explain what happens:



8b Explain how you can digitally sign and send a message m that is allowed to be sent as plaintext. 5pt

You should come up with the following figure, and explain what happens:



Final grade: (1) Add, per part, the total points. (2) Let T denote the total points for the midterm exam ($0 \leq T \leq 45$); $D1$ the total points for part I; $D2$ the total points for part II. The final number of points E is equal to $\max\{T, D1\} + D2 + 10$.