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 What is the difference between a *network service* and a *network protocol*? Is it possible to have different implementations of the same protocol? *5pt*

A service describes what *the network offers to its clients by means of an interface. A protocol is a set of rules describing exactly* how *that service is to be realized through the exchange of messages between computers. These rules may be implemented differently (e.g., using different programming languages), as long as those implementations coincide with the rules.*

1b While considering traditional POTS modems, Bob claims that "...it is unusual to use a 2400 baud modem. Modems usually support 56 kbps." What is wrong with this statement. Motivate your answer. *5pt*

Traditional POTS modems still support 2400 baud, which tells how many times a signal may change per second. The higher data rate is obtained through modulation schemes.

1c What type of multiplexing is used for ADSL? Explain your answer. *5pt*

Frequency division multiplexing: the bandwidth of the cable is partitioned into separate 4 kHz channels. This is typical for FDM.

2a Explain why bit stuffing is needed, and how it works. *5pt*

Bit stuffing is needed when using frame delimiters such as 01111110. In order to send the delimeter as data in a frame, a zero is inserted after five consecutive 1s have been transmitted. The receiver simply removes the zero following 5 1s in row.

2b The physical layer sends data as a "continuous" stream of bits. The data link layer uses chunks of bits, called *frames*. Why switch to frames? *5pt*

Continuous bit streams make it impossible to detect errors, to synchronize transmission rates, etc.

2c Data link protocols almost always put the checksum in a trailer, rather than in a header. Why? *5pt*

Yes, it's a question from the problems section, and it has been asked before. The CRC is computed during transmission and appended to the output stream as soon as the last bit goes out onto the wire. If the CRC were in the header, it would be necessary to make a pass over the frame to compute the CRC before transmitting. This would require each byte being handled twice-once for checksumming and once for transmitting. Using the trailer cuts the work in half.

3a Explain the difference between *go-back-N* and *selective repeat*. *5pt*

With go-back-N, whenever a receiver detects a missing or damaged frame k, all frames succeeding k are ignored and the sender is forced to start retransmission at frame k. With selective repeat, it is possible to retransmit only frame k; the receiver buffers successive frames.

3b Let W denote the receiver's window, and M the maximal sequence number, explain why $W \leq \frac{1}{2}M$. *5pt*

Choosing W and M like this ensures that numbers of outstanding frames (as perceived by the sender) can never overlap with the receiver's previous window of acknowledged frames. In this way, whenever an ACK is lost, we avoid that a retransmission by the sender is interpreted as new *frame by the receiver.*

3c Why is it useful to have a large window size when dealing with slow connections such as with satellites? *5pt*

Having a large window size allows to have many outstanding frames, thus increasing the utilization of the connection.

Part II

4c When a router receives an IPv4 packet, it lowers the packet's time-to-live (TTL) value by one. If the TTL value drops to zero, the packet is discarded and an ICMP message is returned to the sender. How can this mechanism be used to trace the route a packet will follow when going from A to B? *5pt*

This principle is applied in the traceroute *program. It sends packets from A to B with increasing TTL values, starting with 1. Every time it gets an ICMP message back, it will have discovered a router. When it finally has the TTL value set such that B is reached, it will (generally) have discovered a route from A to B.*

4b Peer-to-peer routing is generally efficient in terms of the number of hops, but not in terms of the distance that is travelled in the underlying network. What is meant by this statement? *5pt*

P2P routing takes place between nodes that can be placed at entirely different locations in the Internet. In other words, two neighbors in a P2P network may actually lay in completely different parts of the Internet, effectively requiring long-distance Internet communication.

5a The difference between UDP and IP packets is minimal. Why shouldn't applications use IP directly? *5pt*

It is a matter of separating concerns: the transport layer should offer an interface to applications that hides the underlying network. For connectionless services, it does so by means of UDP. Furthermore, note that IP itself has no notion of ports. This type of address belongs to the abstraction provided by the transport layer and of which applications make heavy use. There is, in general, no way that an application can get an IP packet sent directly to an application.

5b Running TCP across wireless networks may lead to performance problems. Why? *5pt*

The problem is that the underlying network may drop packets due to bad transmission conditions. However, TCP assumes that packet loss is caused by congestion and as a consequence starts to decrease the congestion window size, in turn leading to sending off smaller segments. What it should have done is adopt a more aggressive form of retransmission (perhaps with smaller segments) instead of just merely lowering the segment size.

5c The figure below shows the state-transition diagram for setting up and tearing down a TCP connection. The thick dashed line represents the normal path for a server; the thick solid line that of a client. What happens according to this diagram when the ACK sent by the client when changing state from "SYN SENT" to "ESTABLISHED," is lost? What happens normally in TCP? *5pt*

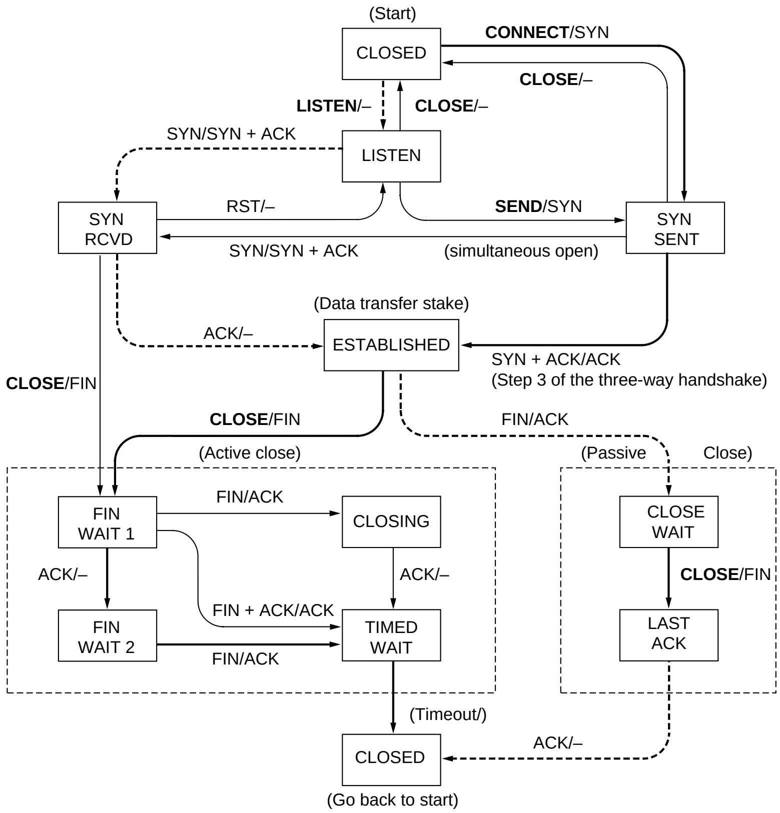

The server blocks in state "SYN RCVD." Normally, a timeout occurs and the server resends the "SYN+ACK" message after which the client can resend its ACK. If even that fails, then both will eventually start closing the connection again.

6a What is the role of DNS when sending mail to *steen@cs.vu.nl*? Be precise! *5pt*

The mail server (or even client) will need to look up the mail server in the cs.vu.nl *domain. It does so through DNS by requesting for the MX record of* cs.vu.nl*. This record contains the DNS name of the mail server, which can then be resolved to the mail server's IP address.*

6b To offload servers, DNS makes extensive use of *caches*. When will DNS caching *not* be effective? *5pt*

When the mapping of name-to-IP address changes a lot, which may be the case with mobile computers.

7a Explain how authentication with public/private key pairs can be established. *5pt*

Let E_A denote A's public key. One approach is that A first sends $E_B[R_A]$ to B, after which B returns $E_A[R_A, S_{AB}, R_B)]$. R_A is a number generated by A and S_{AB} is a shared key generated by B. In the third step, A returns $S_{AB}[R_B]$.

7b Construct a certificate issued by an authority *CA*, telling that E_A is the public key of *A*. Explain the content of your certificate. *5pt*

$[E_A, A, CA, K_{CA}(E_A, A)]$*, where K_{CA} is the private key of CA.*

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* $\max\{T, D1\} + D2 + 10$.