


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 The following figure shows the relation between hosts and a subnet. What would be an appropriate topology for the subnet when we assume that a large number of LANs together should form a supercomputer? Motivate your answer.

5pt




The problem we need to solve is that routes between the different LANs should be short. One possibility is to build a fully-connected graph, but this may easily lead to many connections. As an alternative, a star topology can be used, provided that the central router can handle all inter-LAN traffic. If the number of LANs grows, hierarchical solutions such as a tree may be needed.

- 1b Expand the figure in (a) so that it reflects an internet.

5pt

One possible extension is by adding several subnets:



- 1c Explain what a backbone is.

5pt

No explicit definition is given in the book, but a backbone is nothing but a computer network that connects several other computer networks. A subnet can act as a backbone. In most cases, there are no hosts connected to the backbone, only special routers or gateways that connect a network to the backbone.

- 2a What is a geosynchronous satellite and what is its main advantage compared to other satellites?

5pt

This is a satellite positioned at approximately 36,000 km above the earth where it rotates at the same speed as the earth itself. For this reason, the satellite appears to be hanging still. The main advantage is that because you always know the position of the satellite, there is no need to "track" it down.

- 2b When it comes to mobility issues, communication through low-orbit satellites and communication through cellular radio are very much the same. Explain. 5pt

With low-orbit satellites, a communicating party will regularly switch between satellites because the spot beam is constantly moving. This will lead to hand-over principles that are, in principle, the same ones that are used to be dealt with cell-switching for mobile users in cellular radio systems. Of course, the actual solutions differ, but the principles are much the same.

- 3a 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. 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.

- 3b Some data link protocols, like SLIP, do not provide any support for error detection or correction. What are the consequences of such a design decision, and how serious are these consequences? 5pt

The consequence is that higher layers will have to find out whether something went wrong. Unfortunately, the next higher-level layer has other work to do, namely forwarding packets to the next best router which by itself may be a resource-intensive task. As an alternative, routers may decide to leave the detection for what it is and let the transport layer or higher layers deal with it. The drawback of that approach is that corrupted data may be needlessly passed through an entire network before being detected as garbage.

- 4a What is the difference between using a switch and using a hub in an Ethernet network? 5pt

A hub is just a simple device that copies each incoming frame to each of its outputs. A switch effectively routes an incoming frame to only those outputs to which the destination station is connected. If necessary, it buffers frames to avoid collisions. The end result is that each Ethernet segment connected to the switch gets to see only local traffic, and external traffic that is destined to one of its stations.

- 4b Consider a network consisting of two identical 802.3 Ethernet segments that are connected by means of a bridge. Under which conditions can the bridge be replaced by a simple repeater? 5pt

It's actually quite simple: Ethernet 802.3 specifies a maximum segment length as well as a maximum number of nodes per segment. Assuming the cables are the same, unless these specifications are not violated, the bridge can be replaced.

Part II

- 5a Explain how reverse path forwarding works and what it establishes. 5pt
Reverse path forwarding is used for efficient broadcast routing in networks that do not make use of spanning trees for computing routes (e.g., as in distance vector routing). The method is simple: when a broadcast packet from source S arrives at a router through interface I , the router will forward that packet only if it normally uses interface I to route packets to S .
- 5b Explain how fragmentation in IP works, and how fragmented packets are reassembled at the destination. 5pt
The principle is explained in Section 5.4.6. For IP, it is worth noting that fragmentation is always on 8-byte boundaries, except for the last fragment. A single bit (MF) is used to indicate whether or not we are dealing with the last fragment.
- 5c Some IP packets should (preferably) never be fragmented. Give two examples of such packets. Motivate your answer. 5pt
Packets for which the destination cannot do the reassembly, such as BOOTP. Likewise, ICMP ECHO packets are good candidates for keeping together in order to accurately measure the round-trip delay to a node.
- 5d Explain and motivate the use of extension headers in IPv6. 5pt
The main reason for using extension headers is simplicity. By keeping the main header as simple as possible and treating anything extra as exceptional, routers will be able to quickly process packets and do just what they are good at: forwarding packets. An extension header is used to indicate the special things that need to be done, such as handling fragments or instructing routers.
- 6a 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.
- 6b Explain how TCP congestion control works. 10pt
See Section 6.4.6. The keywords that should be mentioned are congestion window and the notion of a slow start.
- 7a Explain how the name `www.cs.vu.nl` is resolved in DNS, given that it is an alias for `soling.cs.vu.nl`, which has IP address 130.37.24.11. Ignore caching. 5pt
Ignoring caching issues, the normal procedure is to look up `www.cs.vu.nl`, which returns the IP address of the name server for the `vu` domain. This server is then requested to look up the server for `www.cs`, and returns the IP address of the name server for `www`, which will then return the name `soling.cs.vu.nl`. Note that the name server for the `cs` domain may already respond with the answer for `www.cs`.
- 7b 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.

Final grade: (1) Add, per part, the total points. Always exclude bonus points. (2) Let T denote the total points for the midterm exam; $A1$ the total points for part I of the first exam; $A2$ the total points for part II of the first exam; $B1$ the total points for part I of this exam; $B2$ the total points for part II of this exam; The final number of points E is equal to $\max\{T, A1, B1\} + \max\{A2, B2\} + 10$.