ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich
Department of Computer Science

Last name

Foundation of Computer Science Information Systems

Profs. G. Alonso, D. Kossmann, F. Mattern, M. Norrie, T. Roscoe, N. Tatbul
May, 25 2007

Last Name : _____

First Name : _____

Leginr. : _____

Rules

- You have 120 minutes for the exam.
- Please write your name and Legi-ID number on all sheets of paper.
- Please write your answers on the exam sheet. Please also use the reverse sides of the exam sheets. If you need more paper, please raise your hand so that we can provide you with additional paper. Write your name and Legi-ID number on those extra sheets of paper.
- Write as clearly as possible and cross out everything that you do not consider to be part of your solution. You must give your answers in English.
- The exam consists of 10 questions (9 two-sided pages). The maximum number of points that can be achieved is 120.

Statement

- If you wish us to publish your results (grade) on a web page, then please sign the following statement.

Signature : _____

1 Dijkstra's Algorithm (5 Points)

Consider the network shown in Figure 1. Use Table 2 to describe all the steps of the Algorithm. Use Table 3 to describe the state of the routing table of node A once the Algorithm has terminated.

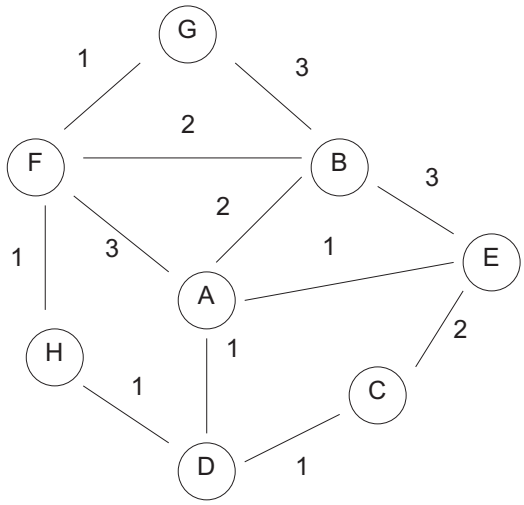

Figure 1: Network

Nodes Visited	Nodes To be visited
A	B, D, E, F

Figure 2: Steps of the algorithm

Node	Next Hop
B	
C	
D	
E	
F	
G	
H	

Figure 3: Routing table

2 CSMA/CD (5 Points)

Consider the linear network shown in Figure 4. Assume that the propagation speed is $2 \cdot 10^8$m/s and the transmission rate of the network is 10Mbit/s. What is the minimum frame size (in bits) necessary to ensure that CSMA/CD will be able to detect all possible collisions on the network?

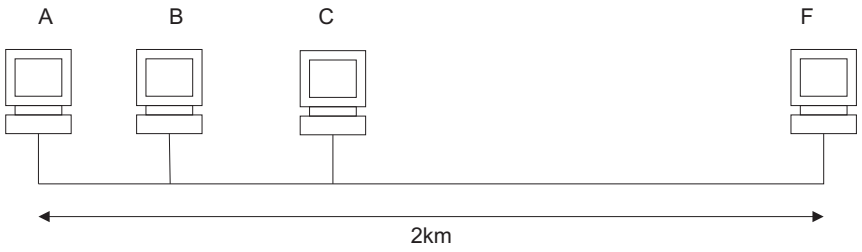

Figure 4: Network

Explain your calculations:

Result: .....

3 Network Utilization (6 Points)

Consider the network shown in Figure 5 with four computers A, B, C and D.

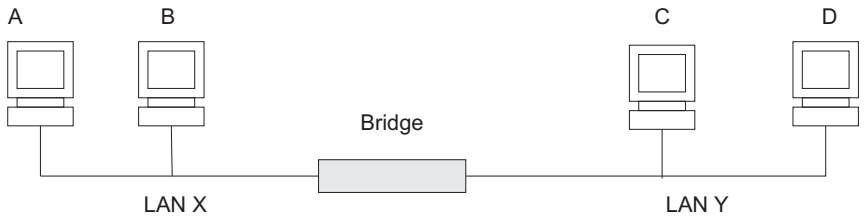

Figure 5: Network

Computer A sends a graphic file of size 10 MB simultaneously to computers B, C and D using UDP unicast packets. Calculate the utilization (the fraction of available bandwidth used) of LAN X, given that each frame carries 1024 bytes of UDP payload data, and transmission is at 50 packets per second to each destination. Assume a MAC header of 8 bytes, an IP header of 20 bytes and an UDP header of 8 bytes.

What is the utilisation on LAN Y?

Calculate the utilisation for LAN Y when the file is sent using multicast packets instead of unicast packets.

4 Sliding Window (5 Points)

Consider a sliding window protocol operating at the data link layer between two stations 1000 km apart, directly connected by an error-free 1.0 Mbit/s link. The frame size used is 500 bits, of which 50 bits are header and 450 bits are data. ACK frames are 100 bits, and the frame processing time is 100 microseconds. Assume that the signal propagation delay is 10 usec/km. Compute the efficiency of the channel under such a setting.

Explain your calculations:

Result: .....

5 Fragmentation (8 Points)

Consider a simple network with 3 nodes arranged in a chain as shown in Figure 6.

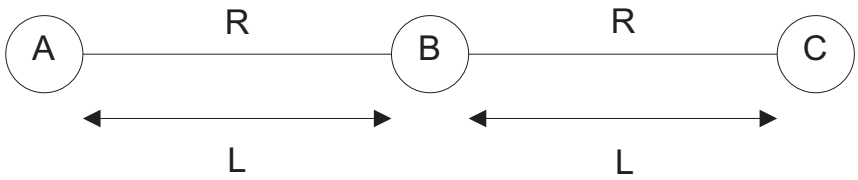

Figure 6: Network

The bandwidth of each link is R Mbit/s. The propagation delay of the link is S and the length of the link is L meters. 12 bits have to sent from node A through node B to node C. Suppose the packet header is 2 bits. And the queuing and processing delays are negligible. What will the delay be to transmit all the 12 bits of data, if they are transmitted using one single packet?

Result: .....

What will the delay be to transmit all the 12 bits of data, if the transmission is split up in two packets of 6 bits each?

Result: .....

What is optimal number of fragment in terms of delay to divide the data into?

Result: .....

6 Bandwidth Allocation (6 Points)

Consider a link of R Kbit/s carrying eight ongoing client-server applications, with each of the application using one TCP connection. If a new application X comes along and opens one new TCP connection what transmission rate will it get?

Result: .....

If the application wants to acquire a rate of R/2, how many TCP connections does it have to open?

Result: .....

Following on the previous question, if the original 8 client-server applications with one TCP connection each all leave the link and 8 new UDP based applications then join the link, would application X be able to keep a rate of R/2?

7 Multiple Choice (25 Points)

Please answer the following questions. For each question, mark the correct answer. There is exactly one correct answer per question. You are awarded 1 point for each correctly answered question; **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.

7.1 ISO/OSI reference model and general questions

1. Which modifications are typically not done by a router?

   ◯ Modification of the IP checksum
   ◯ Decreasing the IP header time-to-live (TTL)
   ◯ Dividing the IP packet into multiple fragments
   ◯ Merging multiple IP fragments to one single packet

2. Suppose a computer with a built-in network interface is disconnected from a network in one building, moved to another building, and connected to the same subnet in the new location. Which of the following statement is true?

   ◯ The computer's MAC address does change while the IP address does not
   ◯ The computer's MAC address and the IP address change
   ◯ The computer's MAC address does not change while the IP address changes
   ◯ Neither the computer's MAC address nor its IP address change

3. Which of the following protocols is typically initiated by a broadcast message

   ◯ HTTP
   ◯ DHCP
   ◯ SMTP
   ◯ SFTP

4. Assume two networks which are connected by a bridge. The only two computers attached are C (client) and S (server) which are located on the same side of the bridge. A client on C is accessing a server application on S using an IP-based protocol. Now S is moved to the other side of the bridge and C is left where it is. Assuming that S is not sending any data, but it is waiting for requests, which of the following statements is true after moving S?

   ◯ The client can immediately successfully open a TCP connection to the server.
   ◯ The client can immediately successfully send a UDP packet to the server.
   ◯ The client application must wait until the ARP (address resolution protocol) cache of C does not contain an entry for S before it can successfully send a UDP packet to the server.
   ◯ The network and/or the computers need to be reconfigured before the client can reach the server.

7.2 Medium Access Control

1. For a certain MAC protocol, if the probability of a frame being damaged is p, what is the mean number of transmissions required to send a frame? Assume that acknowledgements are never lost.

   ◯ $\sum_{k=1}^{\infty} k \cdot (1-p) \cdot p^{k-1}$
   ◯ $\sum_{k=1}^{\infty} k \cdot p^{k-1}$
   ◯ $\sum_{k=1}^{\infty} k \cdot p^{k}$
   ◯ $\sum_{k=1}^{\infty} k \cdot (1-p)^{k-1}$

2. Which of the following devices run a medium access control protocol when connected to an Ethernet segment?

   ◯ Repeaters
   ◯ Bridges
   ◯ Routers
   ◯ None of the above

3. Why the payload of a Ethernet frame needs to be at least 46 bytes?

   ◯ The transmission of small packets is inefficient
   ◯ Collision detection might not work otherwise
   ◯ Ethernet uses a Cyclic Redundancy Check error detection with the polynomial $x^5 + x^4 + x$. The Ethernet payload has to be 46 at least to make sure that polynomial reminder used as parity bits has at least a size of 2 bits.
   ◯ Header, Payload and Checksum sum up to $64 = 2^6$. Packets of this size can be processed more efficiently than packets with a size not a multiple of 2

4. In comparison with slotted Aloha, classic ALOHA

   ◯ Offers better throughput
   ◯ Is significantly easier to implement
   ◯ Requires a global clock
   ◯ If only one station wants to send data, slotted ALOHA results in higher delays

5. In 802.11, the hidden terminal problem occurs when two stations cannot sense each other's signal, but a third station can receive from both. In 802.11, which of the following helps to solve the hidden terminal problem?

   ◯ Carrier sensing
   ◯ RTS (Request to send)
   ◯ CTS (Clear to send)
   ◯ Other than the above

6. Suppose ETHZ and EPFL networks are connected via a 1Gbit link. An alternative way to exchange data is through a dedicated van, capable of carrying 200 tapes of 10GB each. The van takes 3 hours to travel from Zurich to Lausanne. Let L(v) and L(n) be the latencies of the van and the network link, respectively, and B(v), B(n) the bandwidth of the van and the network, respectively. Which statement is true?

   ◯ $L(v) > L(n) and B(v) > B(n)$
   ◯ $L(v) < L(n) and B(v) > B(n)$
   ◯ $L(v) > L(n) and B(v) < B(n)$
   ◯ $L(v) < L(n) and B(v) < B(n)$

7. Assume a set of codewords used in a link level protocol have a hamming distance of d. What is the maximum number of bit errors which can still be corrected when using such a code.

   ◯ $\frac{d-1}{2}$
   ◯ $\frac{d}{2}$
   ◯ $\frac{d+1}{2}$
   ◯ no answer is correct

7.3 Application layer

1. To transfer a file, HTTP is often preferred over FTP because

   - ◯ HTTP achieves a substantially higher throughput than FTP
   - ◯ HTTP works better with network address translation (NAT)
   - ◯ FTP does not support encryption
   - ◯ FTP does not support authentication

2. Since its invention, the Hypertext Transfer Protocol (HTTP) has been extended with persistent connections and pipelining. Which of the following can FTP best be compared to?

   - ◯ HTTP 1.0 (non-persistent, no pipelining)
   - ◯ HTTP 1.1 with non-persistent connections and pipelining
   - ◯ HTTP 1.1 with persistent connections and no pipelining
   - ◯ HTTP 1.1 with persistent connections and pipelining

7.4 Middleware

1. Which of the following cannot be implemented using dynamic binding in RPC?

   - ◯ Load balancing
   - ◯ Redundant servers
   - ◯ Run time client updates
   - ◯ Increased fault tolerance

2. Serialization of data is the process whereby

   - ◯ Data types are mapped between representations in different programming languages
   - ◯ Data types are transformed into a format suitable for sending them through a network
   - ◯ Data types are adjusted to the underlying CPU architecture
   - ◯ Calls by reference are changed into calls by value

3. When using Remote Procedure call, the client stub can be programmed to keep a persistent record of the call to the remote stub. What can be achieved with this strategy?

   - ◯ Exactly once semantics
   - ◯ At least once semantics
   - ◯ At most once semantics
   - ◯ Transactional semantics

4. Which of the following statements does not hold for Lamport clocks? Let C(X) be the lamport time of event X,

   - ◯ If a happens before b in the same process, $C(a) < C(b)$
   - ◯ If a happens in a different process than b, $C(a) \neq C(b)$
   - ◯ If a and b represent the sending and receiving of a message, $C(a) < C(b)$
   - ◯ None of the above

7.5 Routing/Transport layer

1. What is the advantage of connectionless services over connection-oriented services?

   ◯ Lower delays
   ◯ Ordered message delivery
   ◯ Better error recovery
   ◯ None of the answers is correct

2. For a packet arriving on a UDP socket to be delivered to the correct process, it is sufficient for UDP to consider

   ◯ The IP destination address
   ◯ The IP destination address and the UDP destination port
   ◯ The IP destination address, the UDP destination port, the IP source address
   ◯ The IP destination address, the UDP destination port, the IP source address and the UDP source port

3. Consider the effect of using slow start on a line with a 10msec round-trip time and no congestion. The receive window is 24KB and the maximum segment size is 2KB. How long does it take before the first full window can be sent?

   ◯ 30 msec
   ◯ 40 msec
   ◯ 50 msec
   ◯ None of them

4. Assume a channel with bit errors and no packet loss, and a channel with packet loss and no bit errors. Each channel employs a protocol to achieve reliable data transfer. Which techniques do the two protocols have in common?

   ◯ Acknowledgements, retransmissions, sequence numbers
   ◯ Acknowledgements, retransmissions, timeout
   ◯ Retransmissions, sequence numbers, timeout
   ◯ Retransmissions, sequence numbers

5. Inside a network address translation (NAT) router, when should the routing decision (where to forward the datagram) be made in relation to altering the datagram?

   ◯ Before the routing decision
   ◯ After the routing decision
   ◯ Depends on whether the source or destination address is altered
   ◯ The order does not matter

6. Consider a Go-Back-N protocol, where sequences are in the range [0,k-1]. What is the highest allowable sender window size to avoid problems with sequence numbers?

   ◯ $\frac{k}{2}$
   ◯ $k-1$
   ◯ k
   ◯ $2 \cdot k$

7. Consider a Selective Repeat protocol, where sequences are in the range [0,k-1]. What is the highest allowable sender window size to avoid problems with sequence numbers?

   ◯ $\frac{k}{2}$
   ◯ $k - 1$
   ◯ k
   ◯ $2 \cdot k$

8. Assume packets arriving at a router as a Poisson process with mean rate 900 packets/sec. The router takes on average 1msec to process each packet, requiring an exponentially distributed service time. What is the average queue length, and what is the average time a packet waits in the queue?

   ◯ Tqueue = 0 msec, Qlen = 0 packets
   ◯ Tqueue = 9 msec, Qlen = 8.1 packet
   ◯ Tqueue = 10 msec, Qlen = 9 packets
   ◯ Tqueue = 18 sec, Qlen = 16.2 packets

8 Entity Relationship, Relational Data Model (15 points)

Consider the ER diagram shown in Figure 7. The two "1"s at the edges of Relation R indicate that R is a one-to-one relationship.

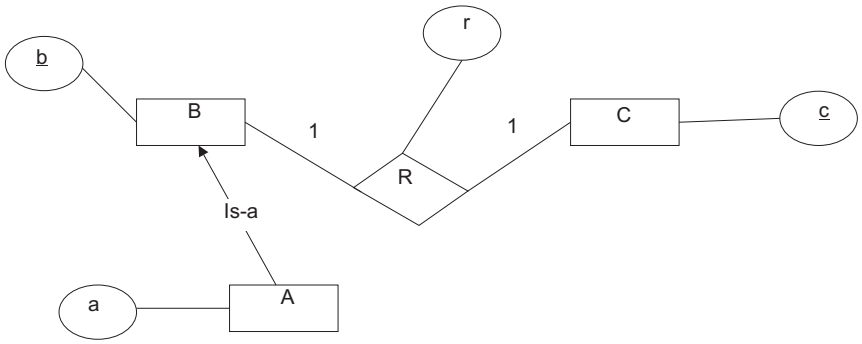

Figure 7: ER diagram

a) Give an SQL schema that implements this ER model. Use the "textbook" rules to translate ER models into a relational model. (The type of all attributes is integer.)

b.) Specify primary key, foreign key, and unique constraints in your SQL schema of (a).

c.) In which (highest) normal form is each relation of your SQL schema? If a relation is not in 4NF, decompose that relation so that the final schema is in 4NF.

9 Query Languages (25 points)

Consider a link table with the following schema: link(source, destination, cost) The link table stores the cost of sending messages from a source node to a destination node. An example is given below:

Source	Destination	Cost
A	B	25
A	C	37
B	A	20
...	...	...

Figure 8: Link table

Note that the communication links are not symmetric.
a.) Give a SQL query that computes all nodes which have a direct link to all other nodes.

b.) Give a SQL statement which deletes all nodes which have a direct link to Node C.

c.) Give a SQL query which computes the average cost of all routes from Node A to Node B with at most one intermediary node on the route.

d.) Consider that the table is represented in XML as follows:

```
<table>
    <tuple source="A", destination="B", cost="25"/> ...
</table>
```

Give an XQuery program that computes all nodes which have a direct link to all other nodes (same as a.).

10 Design of Relational Schemas (20 points)

Consider the following relation:

```
Student(name, hobbies, classes, professor)
```

This relation captures for each student, the name of the student, the student's hobbies (possibly several such as football, tennis), the classes the student attends (possibly several such as compilers, operating systems), and the professors who potentially teach those classes (possibly several such as Gutknecht and Roscoe for operating systems).

a) List all non-trivial functional and multi-valued dependencies of the Student relation.

b) Is the Student relation in BCNF (Boyce Coddd Normal Form)? If yes, explain why. If not, decompose the Student relation so that all relations of the resulting schema are in BCNF.

c) Is the Student relation in 4NF? If yes, explain why. If not, decompose the Student relation so that all relations of the resulting schema are in 4NF.