FREIE UNIVERSITÄT BOZEN
LIBERA UNIVERSITÀ DI BOLZANO
FREE UNIVERSITY OF BOZEN · BOLZANO

Fakultät für Informatik | Facoltà di Scienze e tecnologie informatiche | Faculty of Computer Science

Distributed Systems

Mock Exam

12.6.2015

FIRST NAME		**LAST NAME**	
STUDENT NUMBER		**SIGNATURE**	

Instructions:

You are allowed to bring one handwritten A4 sheet of paper to the exam.

Write First Name, Last Name, Student Number and Signature where indicated. If not, the examination can not be marked.

Do not speak to any other student during the examination. If you speak to another student, your examination will be cancelled.

It is ok to make additional assumptions, but write them down.

Write neatly and clearly.

The real exam will be ~33% smaller. There will be 5 topics, and they will generally be smaller than presented here (e.g. if Coordination was on the exam, then only either Lamport's clocks or the Bully algorithm would be there).

For the 4CP version, only topics 1-4 are relevant.

1. Protocol Stack

(i). Consider a DNS package containing 800 Bytes of data, which is transmitted over an Ethernet with a maximum transmission unit (MTU) of 512 Bytes. Draw the packages that are transmitted over the physical medium, clearly marking the location of headers and data.

(4 points)

(ii) Which service does the TCP protocol provide to protocols in the transport layer above, and which services of the underlying IP layer does it use?

(3 points)

2. Data Link Layer

(i) The following sequences are code words in the (7,4)-Hamming code from the lecture. What are the message bits that the receiver will get after decoding? Did errors occur during the transmission, and if so, where?

(4 points)

1101 000 0011001 0111101

(ii) Consider a link layer in which a bit gets distorted with probability 10^{-6}, and a frame size of 500 bit. In terms of bandwidth, is it more efficient to correct errors with a Hamming code, or to use an error detecting code and retransmit garbled frames? What else than bandwidth could be an argument in favor of one of the two methods?

(4 points)

(iii) The following character encoding is used in a data link protocol:
A: 01000111; B: 11100011; FLAG: 01111110; ESC: 11100000
Show the bit sequence transmitted (in binary) for the four character frame A B ESC FLAG when each of the following framing methods is used:

1. Character count
2. Flag bytes with byte stuffing
3. Starting and ending flag bytes, with bit stuffing

(6 points)

3. Medium Access Control Sublayer

(i) Compare contention with contention-free techniques for medium access sharing. Given one example of a technique in each category, and mention a scenario in which the one technique fares better than the other, and vice versa.

(5 points)

(ii) Draw the signal on the channel for 4 senders A, B, C and D, which each want to transmit frames of 5 bits length as detailed below, using the CSMA/CD protocol. Assume zero channel delay.
A and C want to send one frame immediately, C decides so at time 7, D so at time 15, A decides to send another frame at time 16. Assume a random number generator generating the following sequence: (0,73; 0,24; 0,91; 0,52; 0,48; 0,11; 0,87; 0,02; 0,30; 0,55; …)

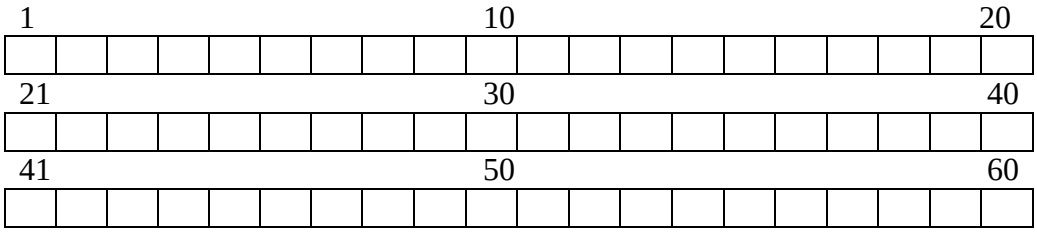

(5 points)

4. Transport Layer

(i) Explain the difference between TCP and UDP. Given one example each for a protocol implemented on top of them, and explain why.

(4 points)

(ii) How does the TCP handshake work, and why is it required?

(4 points)

(iii) Explain how flow control is implemented in the TCP protocol, and simulate 5 rounds of flow control for two network flows A and B, starting at 70% and 0% of the network utilization, respectively.

(6 points)

(iv) TCP uses AIMD (additive increase, multiplicative decrease) for flow control. Could MIAD or MIMD be used as well? Explain your answer.

(4 points)

(v) What are the token bucket and the leaky bucket algorithm used for, and what is the difference?

(4 points)

5. Cryptography

(i) In Cryptography, there is a distinction between symmetric and asymmetric encryption. What is the difference, and what are the advantages and usage scenarios of both methods?

(4 points)

(ii) How can asymmetric keys be used for encryption/decryption versus message signing?

(2 points)

(iii) What are certificate chains and certificate authorities, and what are they used for?

(2 points)

(iv) Encode the plaintext symbols 3, 5 and 10 using the private key (5,13)

(6 points)

6. Coordination

(i) Consider three distributed processes P1 , P2 , and P3 . The processes are involved in the events a, b, . . . , k listed below, which happen at specific points in time, specified as "wall-clock time" (WCT) and measured in ms:

At 0ms WCT:
- a: P2 sends message m1 to P1
- b: P3 reads a startup file

At 100ms WCT:
- c: P1 opens a file containing a user profile
- d: P3 sends message m2 to P2

At 200ms WCT:
- e: P1 receives message m1
- f : P2 receives message m2

At 300ms WCT:
- g: P1 sends message m3 to P2 and P3

At 400ms WCT:
- h: P2 receives message m3

At 500ms WCT:
- i: P2 sends message m4 to P1

At 600ms WCT:
- j : P1 receives message m4

At 600ms WCT:
- k: P3 receives message m3

a. Which of these events are related by Lamport's "happened before" relation? Draw a directed graph where the events are vertices and where there is an edge from E1 to E2 if E1 happened before E2.

(6 points)

(ii). Associate to each event a logical timestamp according to the logical clock algorithm. Use process IDs to break possible ties. What is the linear order of events induced by these timestamps?

(4 points)

(iii) Sketch the Bully Algorithm. Remember there are 3 types of messages:
– election, vote, coordinator.

(4 points)

(iv) Execute the bully algorithm in a network with 6 peers, numbered from 1 to 6, where Peer 2 is the first to notice the crash of the leader number 6.

(6 points)

(v) What happens if two processes notice at the same time that the leader has crashed?

(2 points)

Some more possible exam topics

- **Synchronization**: You could be given some erroneous code, you would have to tell what is wrong, and how to fix it.
- **Link state and distance vector routing**: You could be given a graph/a set of distance vectors, and would need to compute the shortest path or an updated distance vector.
- **IP addresses**: You could be given a block of addresses, and would be asked to split it into subnets of some size. Or could be asked which IP addresses would be directed where given a router that uses longest-matching-prefix routing.
- **Chord**: You could be asked to fill finger tables of nodes, and/or to perform a lookup operation.
- **RMI**: You could be asked to explain the semantics of parameter passing in Java RMI. And how parameters from the calling thread are passed to the called method and vice versa.