TDDD93/TEN2 – Large-scale distributed systems and networks
Final Examination: 14:00-18:00, Thursday, Oct. 22, 2015
Time: 240 minutes
Total Marks: 40
Grade Requirements: Three (20/40); four (28/40); and five (36/40).
Assistance: None (closed book, closed notes, and no electronics)
Instructor: Niklas Carlsson

Instructions:

- Read all instructions carefully (including these)!!!! Some questions have multiple tasks/parts. Please make sure to address *all* of these.
- The total possible marks granted for each question are given in parentheses. The entire test will be graded out of 40. This gives you 10 marks per hour, or six minutes per mark, plan your time accordingly.
- This examination consists of a total of 12 questions. Check to ensure that this exam is complete.
- When applicable, please explain how you derived your answers. Your final answers should be clearly stated.
- Write answers legibly; no marks will be given for answers that cannot be read easily.
- Where a discourse or discussion is called for, be concise and precise.
- If necessary, state any assumptions you made in answering a question. However, remember to read the instructions for each question carefully and answer the questions as precisely as possible. Solving the *wrong* question may result in deductions! It is better to solve the *right* question incorrectly, than the *wrong* question correctly.
- Please write your AID number, exam code, page numbers (even if the questions indicate numbers as well), etc. at the top/header of each page. (This ensures that marks always can be accredited to the correct individual, while ensuring that the exam is anonymous.)
- Please answer in English to largest possible extent, and try to use Swedish or ”Swenglish” only as needed to support your answers.
- If needed, feel free to bring a dictionary from an official publisher. Hardcopy, not electronic!! Also, your dictionary is not allowed to contain any notes; only the printed text by the publisher.
- Good luck with the exam.

Part A: Distributed Systems

Question 1 (4 points)

Transparency plays a central role in some distributed systems. Consider a simple multi-tier system with three levels: a user interface, an application server, and two replicated database servers. Assume these layers are implemented as a distributed cloud service at different geographic locations and that the average round trip time (RTT) between the machines used in the consecutive layers (starting with the top-tier layer) is 30ms and 20ms, respectively. Consider a workload (set of calls) with two different types of "jobs" (call types). The first type results in fully synchronized calls in which the application server requires 50ms total processing and the database requires 100ms processing to satisfy the request. The second type does not require any database access, is fully synchronized and requires 50ms processing at the application server.

(a) For each of the two types of jobs, how much time is the client process locked from the moment it makes the request to the application server? You can assume that no large data is transferred between the layers such that the call and responses fits within a single package, and that messages do not need to be acknowledged. Please explain your answer and illustrate with a figure.

(b) Assuming 50% of the clients make each type of requests, what is the average response time (assuming no competing jobs or other reasons for queuing).

Please give concrete examples of two types of transparency that are provided in the above example. Remember to explain your answers.

Question 2 (4 points)

Assume that you have three processes p1, p2, and p3 which are implementing Lamport's clocks. There are many events that take place at these processes, including some messages being sent between the processes. In the figure below we use circles and arrows to specify in-processor events and messages being sent between processes, respectively. Please provide the logical timestamps associated with each event. You can assume that all three clocks start at zero, at the left-most point in time. (Also, explain how the processes would adjust their clocks if using Lamport's logical clocks.)

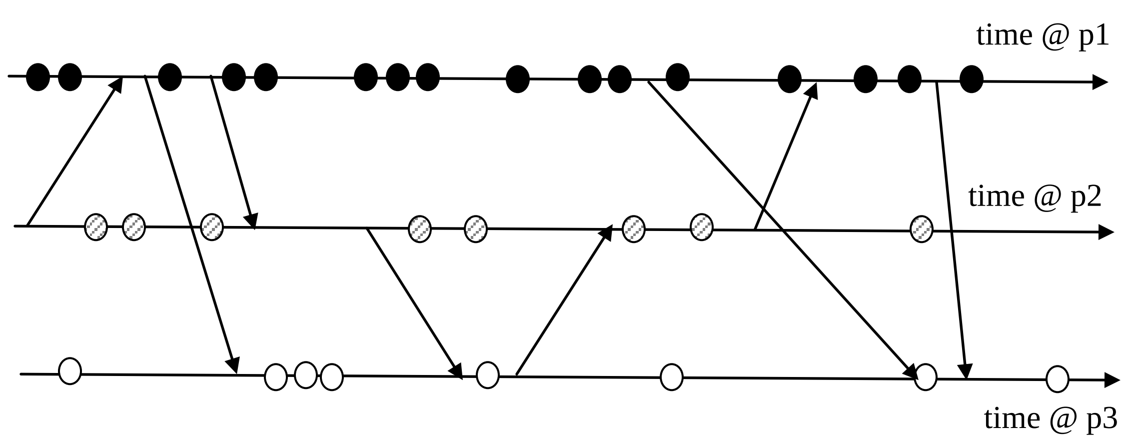

Question 3 (2 points)

In the context of remote procedure call (RPC), please describe and compare at least two potential actions that a server orphan can take after the client has crashed while the server was computing.

Part B: Methodology

Question 4 (4 points)

When designing experiments, it is important to carefully identify the most appropriate factors, levels, and metrics to consider. Consider a researcher wanting to assess the performance of a webserver. The researcher has identified three factors of interest: (i) the request rate, (ii) the job size, and (iii) the processor speed. For each of these factors, the researcher has identified 8 levels (each) of interest, including identified a default request rate, job size, and processor speed. Let us call the request rate levels R1, R2, ..., R8; the job size levels S1, S2, ..., S8; and the processor speed levels P1, P2, …, P8. Please estimate the number of experiments that the researcher would need to perform if performing (a) one factor experiments with the default scenario as baseline, (b) two factor experiments with the default scenario as baseline, and (c) full factor experiments. Also, please explain which experiments would be performed in each case.

Question 5 (3 points)

Consider a system with two states: "on" and "off". Assume that the system is "on" whenever there are jobs to serve and the system instantaneously can go between the "on" and "off" states whenever a new job arrive to an empty system or the system is done serving all jobs, respectively. Furthermore, assume that the system only can serve one job at a time (as with any G/G/1 queue system), on average 100 jobs/second arrive to the system, each job on average takes 20ms to serve, and each job stays in the system for on average 40ms.

a) How many jobs are on average in the system?
b) Assuming that the "on" state consumes 100 Watt and the "off" state 10 Watt. What is the average power consumption of the system, given the described workload and system characteristics?

Question 5 (3 points)

Describe the main difference between active measurements and passive measurements. Explain which approach introduce the most traffic overhead and use a concrete example to explain how these techniques can be used to build a map of the Internet topology. In your explanation, please use existing protocols/techniques and how they can be used for active and passive measurements.

Part C: Multicore and Parallel Programming

Question 7 (2 points)

Questions on parallel computer architecture concepts

a) Explain the concept of "Hardware multithreading". (1 point)
b) Which interconnection network topology is most suitable for realizing scalable on-chip networks connecting cores in homogeneous many-core processors? Motivate your answer. (1 point)

Question 8 (4.5 points)

Questions on thread programming

a) Given a very large string of N characters stored in shared memory. Write a shared-memory parallel program using threads (pseudocode is fine, explain your code) that counts the total number of occurrences of the substring "OK" in the string. Make sure that your program is free of race conditions, and try to achieve cache-friendly memory access patterns. Explain your solution. (3 points)
b) Derive the asymptotic worst-case parallel execution time, parallel work, and parallel cost for your algorithm as functions (big-O notation) in N and the number P of processors, using the (EREW) PRAM model. (You may assume that $P << N$.) (1.5 points)

Question 9 (3.5 points)

Questions on Design and Analysis of Parallel Algorithms

a) How does the PRAM model relate to the traditional sequential RAM (Random Access Machine) computation model? Which parameter(s) does the PRAM cost model expose? What properties of a real parallel computer does it abstract from? Why can it be beneficial to nevertheless start with developing a PRAM algorithm when developing an algorithm to be run on a concrete parallel computer? (2.5 point)
b) What is the "critical path" of a parallel computation (for a given input)? What does its length tell us about the parallel execution time that we can expect on a concrete parallel computer? (1.5 point)

Part D: Embedded Systems

Question 10 (4 points)

a) Describe, using a flow graph, the design flow of an embedded systems, from an informal specification to fabrication.
b) Give short comments on the design steps which belong to the system-level.
c) Why is the proposed design flow better than the traditional one?

Question 11 (3 points)
What is an Embedded System? What makes it different from other applications? Why is it difficult to design?

Question 12 (3 points)
Think at the sources of power dissipation as we discussed at the lectures. What are main opportunities to reduce power consumption?

Good luck!!