TDDD93/TEN2 – Large-scale distributed systems and networks
Final Examination: 8:00-12:00, Thursday, Aug. 20, 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 13 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)

Mutual exclusion. Consider a simple scenario in which there are five nodes A, B, C, D, and E. Use a sequence of figures to illustrate and explain the message sequences and coordination between these nodes when (i) node A acts as a central coordinator for a shared memory resources (that all five nodes can use) and both nodes B and C almost at the same time decides that they want to write to the resource, (ii) the bully algorithm is used to elect a new coordinator when node A crash.

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.)

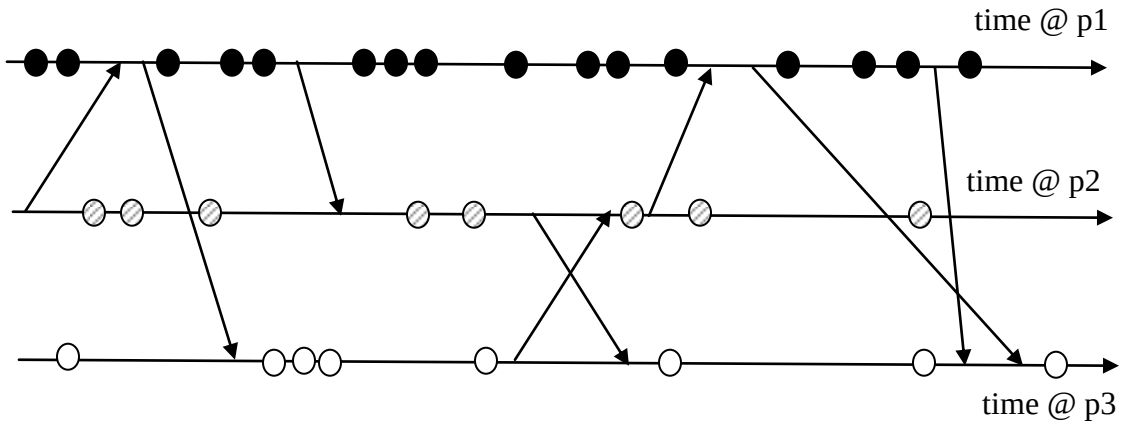

Question 3 (2 points)

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. Within this context and example scenario, please give two concrete examples of two different types of transparency and explain how transparency is used here to provide improve service for the end users.

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 6 (3 points)

Power law distributions (and other related "heavy tailed" distributions) are often observed in nature and large distributed systems.

a) How does "heavy tailed" distributions relate to exponential distributions?
b) Plot both a power-law distribution and an exponential distribution side-by-side on both a lin-lin plot (with both linear x-axis and linear y-axis) and a log-log plot (with both x- and y-axis on log scale).

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) Name and briefly describe one type of interconnection network where the node degree is constant (i.e., independent of the number of nodes). (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 (1 points)

Question on MPI

a) Name and explain one collective communication operation in message passing / MPI programming. (1 point)

Question 10 (2.5 points)

Questions on Design and Analysis of Parallel Algorithms

a) What is the purpose of using a parallel cost model when designing parallel algorithms and programs (as opposed to doing the analysis directly for the concrete target architecture that we have at hand for program execution)? (1 point)

b) Give an example scenario and explanation of a parallel speedup anomaly. (1.5 point)

Part D: Embedded Systems

Question 11 (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 12 (3 points)

What is an Embedded System? What makes it different from other applications? Why is it difficult to design?

Question 13 (3 points)

Why is power consumption an important issue in today's computer systems?

Good luck!!