TDDE35/TEN2 and TDDD93/TEN2 – Large-scale distributed systems and networks
Final Examination: 8:00-12:00, Wednesday, Aug. 22, 2018
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 and utilize figures and tables to the largest extent.
- 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.

TDDE35/TEN2 and TDDD93/TEN2
Final Exam: 8:00-12:00, Wednesday, Aug. 22, 2018

Part A: Distributed Systems

Question 1 (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.)

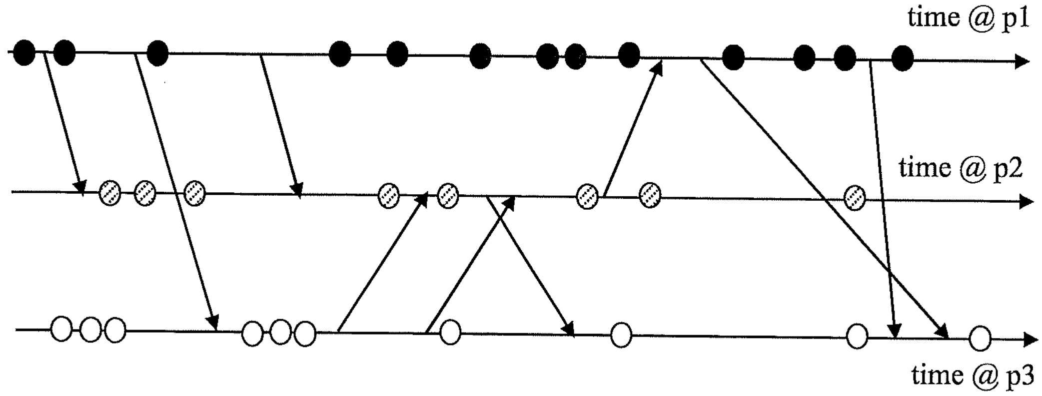

Question 2 (4 points)

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 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. You can assume that each write access (to memory) takes 1 second and that there is 10ms between the times when B's and C's decisions, and that the round trip times between the nodes are random in the approximate range 60-120ms.

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 (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.

Question 5 (3 points)

You have performed large-scale measurements and are now using scatter plots (i.e., you plot each data point individually in the x-y plane) to visualize your results. When visualizing the results you notice clear trends.

(a) What does it mean if all points end up being on a straight line on a lin-lin plot (i.e., both axes on linear scale)? Show, explain, and try to provide example equations to interpret the results.
(b) What does it mean if all points end up being on a straight line on a log-log plot (i.e., both axes on logarithmic scale)? Show, explain, and try to provide example equations.
(c) What does it mean if all points end up being on a straight line on a lin-log plot (i.e., one axis on linear and the other on logarithmic scale)? Show, explain, and try to provide example equations.

Question 6 (4 points)

Scalability comparison of peer-to-peer and client server. Please derive expressions for how much time it takes to distribute file from one server to N clients/peers when using the client-server model and peer-to-peer, respectively. You can assume that the server has an upload bandwidth U, that each peer i has an upload bandwidth of u_i, that each peer i has a download bandwidth d_i, and that the file is of size F. Then use the expressions to and plot the distribution time as a function of the number of clients/peers in the system for the two delivery models, for the case when using the normalized file size $F = 1$ (i.e., the file size is measured in units of the file size itself) and $U = u_i = 1$ (i.e., the upload rate is measured in the time units that it takes the server to upload one copy of the file and the clients have the same upload rate) and $d_i = 2$ (i.e., the maximum download rate of a client is twice its upload rate). For simplicity, assume that the file can be broken into infinitesimally small chunks and that a client can start help out as soon as it has obtained such small chunk.

Part C: Multicore and Parallel Programming

Question 7 (3 points)

Questions on parallel computer architecture.

a) Name and shortly describe an interconnection network topology that is suitable for on-chip networks in many-core processors. (1p)
b) How does "hardware multithreading" differ from "multicore" technically (be thorough), and what do they have in common from the software point of view? (1.5p)
c) Why did the transition from single-core to multi-core CPUs not solve the "Memory Wall" problem? (0.5p)

Question 8 (1 points)

Question in thread programming.

a) What is the purpose of a thread pool? (1p)

Question 9 (4 points)

Questions on MPI/algorithm design.

Given is a cluster computer with P (P>1) nodes. Assume the input (a huge string of size N) is already distributed and each node has its part of the overall input, stored in a large array A of N/P characters in its main memory.

i) Write a message-passing parallel program (MPI or pseudocode is fine, explain your code) that uses all P nodes in parallel to count the total number of occurrences of the character 'X' in all nodes' arrays A. The total number should finally be known to each node, i.e., stored in a variable in its memory. Be thorough, and explain your code carefully. (2p)
ii) Make sure to explain all communication operations in (i). Draw a figure for P=4 (fully connected) nodes showing the distribution of A and the flow of messages between the nodes. Which of these communication operations are point-to-point communications and which ones are collective communication operations? (1p)
iii) Derive the asymptotic worst-case parallel execution time for your algorithm as an expression (big-O notation) in N and the number P of nodes. (1p)

Remark: If you do not know how to write message passing parallel programs, you could instead solve (i) and (iii) for the shared memory (multithreaded) programming model, though with half the amount of points each because it is much easier.

Question 10 (2 points)

Question on theory.

i) A parallel program fulfilling the assumptions of Amdahl's Law takes 100 seconds if run on one processor and 60 seconds if run on 2 processors. What is the upper bound of the relative speedup that this program could achieve for any number of processors? Explain your answer (calculation). (1p)
ii) When is a parallel algorithm cost-optimal? (0.5p)
iii) Give an example of a parallel algorithm that is not cost-optimal. (0.5p)

TDDE35/TEN2 and TDDD93/TEN2
Final Exam: 8:00-12:00, Wednesday, Aug. 22, 2018

Part D: Embedded Systems

Question 11 (4 points)

Consider an application modelled as the task graph below. Each task, when activated, consumes one message on each input edge and generates, at termination, one message on each output edge. The task graph is executed on the architecture shown in the figure. Execution times of the tasks, when executed on the corresponding processor, are shown in the table. All messages transmitted over the bus, between tasks mapped on different processors, consume 2 time units to reach the destination. Communication between tasks mapped to the same processor is considered to not consume any time.

Propose an efficient task mapping (indicate on which processor each task is executed) and a corresponding static (nonpreemptive) schedule for the application. Illustrate your schedule as a Gantt chart (similar to the way we captured schedules in Lecture 1&2).

Try to achieve a maximum delay (the time interval between the start of T1 and the finishing of T7) of 46 time units!

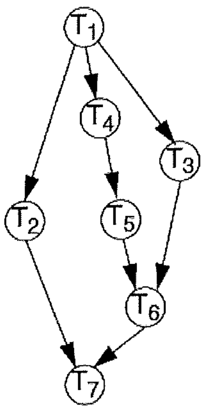

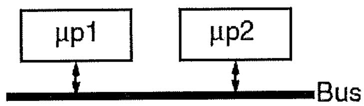

Task	WCET	
	μp1	μp2
T_1	5	6
T_2	12	15
T_3	10	11
T_4	5	6
T_5	3	4
T_6	17	21
T_7	10	14

Question 12 (3 points)

In the lectures we have particularly emphasized three design steps: architecture selection, task mapping, elaboration of a schedule. Explain, in short, what each step is doing. Illustrate the three steps by a small example.

Question 13 (3 points)

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

Good luck!!