

TDA596 / DIT240, 7.5 hp, (2nd academic period 2012/2013)

Exam: Distributed Systems

18 Dec. 2012.

Examiner: Olaf Landsiedel

Contact: olaf@chalmers.se, cell phone: 070-9702443, office phone: 772 17 06

Means allowed: Nothing except paper, pencil, pen and English - xx dictionary.

Please answer questions 1 to 6.

General information: All questions should be answered in English. Write clearly and use the pages in a clever way so it is easy to read. Each question answer should be started on a new paper. All answers should be motivated, explained, elaborated, detailed, precise and accurate.

Important suggestion: Read all questions before answering. Plan your time so that you can (at least) write a brief answer to all questions. Please notice the points that are given to each (sub)question.

Grading: GU: G 24p, VG 48p; CTH: 3:a 24p, 4:a 36p, 5:a 48p of maximum 60 points.

Review: Please keep your exam code. Information about individual exam reviews will be published on the course website.

1. Basics about Distributed Systems (10 points)


- 1 a) (2 points) Define the term “Distributed System” in your own words. Be brief and precise.

- 1 b) (4 points) When designing or developing a Distributed System, its distributed nature presents a number of challenges. List at least 4 challenges.

- 1 c) (4 points) Select two challenges you listed in 1b) and detail on them. Answer the following two questions for each of them: Why do these make developing Distributed Systems challenging? What mechanisms did we study in the course to deal with them? Be brief and precise.


2. Naming (10 points)

- 2 a) (2 points) Define the terms “naming” and “name resolution” in Distributed Systems in your own words. Be brief and precise.
- 2 b) (3 points) In the lecture we discussed the concept of Chord. Chord is a Distributed Hash Table (DHT). Answer the following questions about Chord:
- 2b i. What operations does a DHT, e.g., Chord, provide?
 - 2b ii. In Chord, how many hops does it take on average to lookup a data item? (Assume that the number of nodes in Chord is n).
- 2 c) (5 points) Finger tables in Chord
- 2c i. What is the finger table in Chord used for?
 - 2c ii. Below you find a picture of a Chord ring (on the left) and the finger table of node 80 (on the right). Please complete its finger table, i.e., list to which nodes the fingers point. In this example the finger table size is 7.




3. Clocks and Synchronization (10 points)

- 3 a) (2 points) In the course we discussed the concept of clock synchronization for physical clocks. Why is clock synchronization important in Distributed Systems that rely on physical clocks?
- 3 b) (2 points) Please answer the following questions about logical clocks.
- How do logical clocks differ from physical clocks?
 - When do logical clocks “tick”, i.e., when are they incremented?
- 3 c) (6 points) In the course, we discussed the concept of Vector Clocks. Please answer the following questions about Vector Clocks.
- How many clocks does each node maintain?
 - Upon receiving a message, what does a node do with its vector clock(s)?
 - Below you see a figure of a vector clock for two nodes (p_1, p_2). All vector clocks are initialized to zero. Please list the vector clocks for events A, B, C.



4. Consistency and Replication (10 points)

- 4 a) (2 points) Define the term replication. Be precise and brief.
- 4 b) (3 points) In the lectures we discussed the concept of “Eventual Consistency”. Explain it briefly.
- 4 c) (5 points) We discussed the concept of “Total Ordering” as one consistency model.
- 4c i. Please describe this concept briefly.
- 4c ii. Below you see two figures showing an execution (sending and receiving of messages) on three nodes (P1, P2, P3). Two messages ($m_{1,1}$, $m_{3,1}$) are sent. Please note for each figure whether the depicted execution shows total ordering. Explain your decision; be precise and brief.



5. Fault Tolerance (10 points)

- 5 a) (2 points) List at least two typical failures in Distributed Systems.
- 5 b) (2 points) Define the term “Fault Tolerance”. Be brief and precise.
- 5 c) (6 points) We discussed the “Three Generals Problem” and a solution.
- 5c i. In the “Three Generals Problem” there are honest generals and dishonest generals (traitors). What is the goal of the honest generals? What is the goal of the traitors?
 - 5c ii. In the lecture we introduced an algorithm to enable agreement among the generals. The algorithm has multiple phases. Explain the algorithm and its different phases. Be brief and precise.
 - 5c iii. Under what conditions can the generals achieve agreement: How many *honest* generals are at least required, assuming that there are k dishonest ones?

6. BitTorrent (10 points)

- 6 a) (4 points) Explain the basic concept of BitTorrent. Explain the terms “Seed(er)”, “Leech(er)”, “Tracker”, and “Swarm”,
- 6 b) (6 points) Explain the concept of Tit-for-Tat that BitTorrent uses to avoid “free-riding”, i.e., ensuring that nodes that download also upload. What roles play “choking”, “unchoking” and “opportunistic un-choking” in this context?