521290S Distributed Systems, mid-term exam 2

23.02.2024

University of Oulu

Please read this section carefully!

The exam consists of **10 questions** from 5 topics. You have **90 minutes**, and you can reach 10 points in total. Every question is worth 1 point, and can be graded *incorrect* (0 pts), *partially correct* (0.5 pts), and *correct* (1 pt).

Further remarks:

- **This is a closed book exam.** Put away all lecture material.
- **Make sure that your writing can be read. We cannot grade what we cannot read.**
- **Put away all electronic devices** (mobile phone, laptop, iPad, etc.). If we spot you working on an electronic device during the exam, we will assume that you were looking up lecture material and consider this cheating.
- Please ask if you have trouble understanding a question.
- We are looking for very compact answers. For each question, one to three (meaningful!) sentences are sufficient.
- Do not write tautologies (e.g., *"A distributed system is a system that works in a distributed way"*).
- In general, the exam paper should have enough space for all your answers. If you run out of space, ask us for more paper. You are not allowed to write on your own paper.

Good luck!

---

6 Naming

6.1 In the figure, you can see a simplified Chord system with a 4-bit identifier space. Grey circles denote nodes, while black/dotted circles denote other entities. For the nodes, finger tables are given. Consider starting at the node with ID 1 and resolving key k = 14.

In your answer, provide the IDs of the nodes which are used to resolve k, i.e., which look up the next node in their finger tables, separated by commas. ***Name both the end node and the starting node.*** *For instance, if you think that the name resolution is done by starting at ID 1, then being forwarded to the node with ID 3 and then to the node with ID 6 (which is the end note), then you should provide "1, 3, 6" as your answer. Please stick to this writing style, else, your answer might be graded wrongly.*

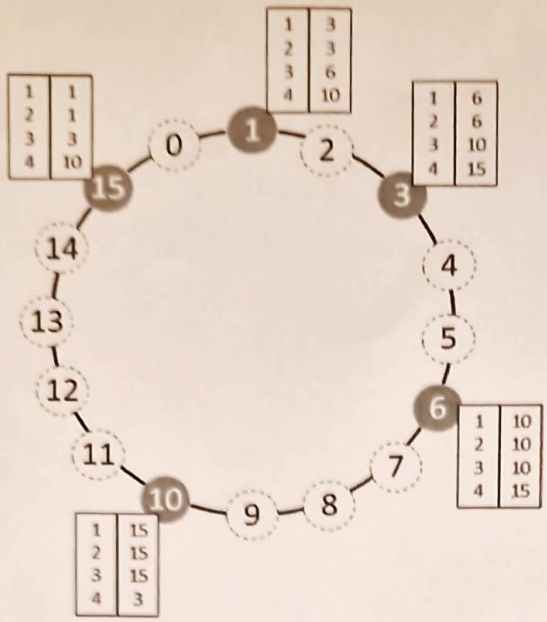

6.2 Below you will find three examples for naming information. Are these either examples for flat naming, structured names, or attribute-based naming?

a) uid=lauri.loven, ou=People, dc=example, dc=fi
b) 54-92-96-4A-4B-55
c) www.oulu.fi

7 Consistency and replication

Consider the two concurrently executing processes P1 and P2 sharing data store x (time increases from left to right for each process individually):

P1: W(x)a _____ **W(x)c**

P2: _____ **R(x)a** _____ **W(x)b** _____

7.1 State and shortly explain all constraints that affect the ordering of operations when the data store x is **sequentially consistent**.

7.2 Show all sequences of read and write operations that are valid given the constraints in question.

8 Fault tolerance

8.1 Explain in one sentence each how large a k-fault tolerant group needs to be in order to tolerate either (i) timing faults and (ii) Byzantine faults. Justify the group sizes.

8.2 Explain the major differences between a Response Failure (also known as Common-Mode Failure) and an Arbitrary Failure (also known as Byzantine Failure) in max. three sentences.

9 Security

9.1 *During the establishment of a secure channel between two parties, after the authentication phase has completed, the communicating parties generally use a unique shared session key for confidentiality, which is discarded when the secure channel is no longer used. Name two benefits of this approach.*

9.2 *Name one reason why we typically digitally sign a message digest produced by a hash function instead of the message itself.*

10 Emerging topics

10.1 *The university has installed surveillance cameras at the entrances of all its buildings, which capture and transmit high-definition video. The video is received by a service that executes a face recognition application, and stores the video and the results of the processing. These data include personally identifiable information, such as faces of students, their identities and other personal data. You have the option to host this service either at the University's on-premise servers or in the Cloud. Which option would you select and why?*

10.2 *An online game for mobile devices involves rendering in real time complex game scenes. This requires a lot of processing and has stringent latency requirements so that a smooth user experience is provided to the player. The application developers have created the game software in such a way that rendering can be offloaded to a remote service component, which sends back game scenes to the mobile device. This involves streaming of real-time high-definition video to the device. Would you run this rendering service at an edge computing server or in the Cloud and why?*

BONUS Question (1 point):

Consider centralized Key Distribution Center (KDC) and Certificate Authority (CA) servers in managing encryption keys. Briefly analyze the advantages and disadvantages of using centralized servers for key management.