



UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

B.Sc. Engineering

2011 Intake Semester 8 Examination

CS4262 DISTRIBUTED SYSTEMS

Time allowed: 2 Hours

February/March 2016

ADDITIONAL MATERIAL: *None*

INSTRUCTIONS TO CANDIDATES:

1. This paper consists of **5** questions in **6** pages.
2. Answer any **4** questions.
3. Start answering each of the main questions on a new page.
4. The maximum attainable mark for each question is given in brackets.
5. This examination accounts for 50% of the module assessment.
6. This is a closed book examination.
NB: It is an offence to be in possession of unauthorized material during the examination.
7. Only calculators approved by the Faculty of Engineering are permitted.
8. Assume reasonable values for any data not given in or with the examination paper. Clearly state such assumptions made on the script.
9. In case of any doubt as to the interpretation of the wording of a question, make suitable assumptions and clearly state them on the script.
10. This paper should be answered only in English.

Question 1 (25 marks)

The CAP theorem states that it is impossible for a distributed system to simultaneously provide Consistency, Availability, and Partition Tolerance.

- (i) Using suitable examples, explain the meaning of Consistency, Availability, and Partition Tolerance in the context of distributed systems. [6]
- (ii) Which property(ies) among Consistency, Availability, and Partition Tolerance is/are affected in the “split brain syndrome”? [3]

Students studying for final exams go out to dinner together every evening (as this their last few days on campus). They study at different locations, so they need a protocol to agree on the dinner destination. They decide to write software that runs on their computers and communicates over the network. Each starts a dinner program on his/her computer, in the morning. Before 6pm each student proposes a dinner location, and then communicates and exchanges his/her location with others. Based on these, they attempt to decide the dinner location. Once a location is decided, it will not change.

- (iii) What type of a network/overlay topology would you recommend for the communication among software run by the students? Justify. [5]
- (iv) What are the challenges that the students’ software needs to overcome to agree on a final dinner destination? [5]
- (v) Outline a high-level solution that could be used to agree on a dinner destination. [6]

Question 2 (25 marks)

- (i) Following is a high-level diagram of a goods dispatch platform (aka. courier service).



Source: <http://srstechnologiesllc.com/mobile-service-platform/>

Customers who wish to use a courier service can contact trucking and transportation companies via the platform and get their goods delivered to desired destinations. For example, suppose Kamal wants to send a box to Rani. Kamal visits the courier service website (Step 1 in figure) and searches for a potential trucking and transportation company in the region. Kamal may also refine his search based on the pricing, delivery options, delivery schedule, and availability of a nearby track or a collection center. Once a company is selected, Kamal's request is forwarded to a driver (via a mobile app, Step 2) who is currently in or will be visiting Kamal's neighborhood soon. The driver receives details such as collection location, drop off location, contact details of Kamal and Rani, and any delivery instructions. The driver then comes and collects the box (Step 3). Then the box is delivered to Rani (Step 4). Finally, the settlement for the provided service is given to the driver's company (Step 5).

- a) What type of a distributed system architectural style would you recommend for the goods dispatch platform? Briefly justify. [4]
- b) Which one would you recommend: **Persistent** or **Transient** communication between the goods dispatch platform and drivers? Justify. [3]
- c) Which one would you recommend: **Asynchronous** or **Synchronous** communication between the goods dispatch platform and drivers? Justify. [3]
- d) Do you recommend **having message queue** at the goods dispatch platform or **not having message queues**? Justify your recommendation. [3]
- e) Using a suitable diagram illustrate your solution based on recommendations for questions (b), (c), and (d) above. [5]
- f) Do you recommend implementing the goods dispatch platform as a **dedicated server(s) in a server room** or **virtual server(s) on the cloud**? Justify your recommendation. [4]
- g) Which one would you recommend: having a **single server** or **multiple servers** for goods dispatch platform? Justify. [3]

Question 3 (25 marks)

- (i) Propose a suitable distributed communication solution for each of the following applications? Justify your suggestions.
 - (a) A taxi dispatch service that not only considers customer requests, but also the location of taxis, real-time traffic, and weather data while deciding what taxi to dispatch. [5]
 - (b) A personal health tracking and recommendation service that tracks every move of a person, his/her eating habits, and then provides recommendations on daily routines, exercises, and meal options. A person's activity and eating habits are to be tracked via smart watches, smart phones, tablets, wireless weight scales, desktops, etc. [5]

- (ii) Suppose a local online retailer (e.g., takas.lk and wow.lk) is planning to generate a customer-specific homepage, where the contents on the homepage is customized based on the past interactions and preferences of the customer. Designers of the website know that detecting and predicting customer preferences require collection of large volumes of data over time and extensive data processing. Moreover, to provide a rich customer experience they are planning to use high-quality graphics and animations.
- (a) Discuss how cloud computing can be used for this requirement. [5]
- (b) What type of a cloud computing solution would you recommend: IaaS, PaaS, SaaS? Justify [5]
- (c) While implementing the Content Delivery Network (CDN) for this website, do you recommend HTTP Redirects based CDN or Domain Name Service (DNS) based CDN? Justify. [5]

Question 4 (25 marks)

Zipf's Law states that in the English language, the frequency of any word is inversely proportional to its rank in the frequency table (i.e., when words are ranked based on their frequency of occurrence). Thus, the most frequent word will occur approximately twice as often as the 2nd most frequent word, three times as often as the 3rd most frequent word, and so on. This can be formally written as the probability of encountering the r -th most common word is roughly $P(r) = 0.1/r$.

Suppose we want to check whether the Zipf's law proposed in 1935 is still valid by counting the frequency distribution of word occurrences of the top 100 books ranked by the New York Times during the last 25 years.

- (i) Is this a suitable problem to be solved using cloud computing? Discuss. [4]
- (ii) What type of a file system do you recommend to store the books? Justify. [3]
- (iii) Outline a solution that shows how you can use Map-Reduce to check whether the Zipf's law is still true. The answer should provide pseudocode for *map* and *reduce* functions. [14]
- (iv) What type of load balancing would you recommend while solving this problem? Justify. [4]

Question 5 (25 marks)

- (ii) Lamport's Timestamps are calculated based on the 2 rules given below.

Processes update their logical clocks and transmit their logical clock values in messages as follows:

LC1

$L_i := L_i + 1$, before each event is recorded at processes P_i


LC2

When P_i sends a message m , logical clock value $t = L_i$, is piggybacked with the message

On receiving (m, t) , a process P_j computes $L_j := \max(L_j, t)$

Then computes $L_j := L_j + 1$ before logically timestamping the event $receive(m)$

- a) Label the following diagram with Lamport's Timestamps.



[4]

- b) Identify 4 event pairs with the same Lamport time stamp. [2]

- c) How can we build a total order of events such that 2 events will not have the same time stamp? What will be the new time stamps for the pairs of events you identified in question b above? [4]

- (ii) Suppose Ms. Kumari is a customer of a multi-national bank. She is currently in the USA on a business trip. Her bank maintains
- three**
- replicas of "customer accounts" database, each in the USA, UK, and Sri Lanka. Kumari is a customer of the Colombo branch. Her bank also has an online banking portal, which connects to respective database depending on the customer's branch.

- (a) What concurrency issues can occur if Ms. Kumari tries to withdraw money from an ATM in the USA, while her branch in Colombo tries to deposit interest to her account? Explain using an example. [4]

- (b) What concurrency issues can occur if Ms. Kumari tries to transfer money via the online banking portal while she is in USA, while her branch in Colombo tries to deposit interest to her account? Explain using an example. [3]

- (c) Using a diagram explain a suitable solution(s) to address issues discussed in (a) and (b). You may recommend established techniques with justification. [8]

----- END OF THE PAPER -----