UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

MSc in Computer Science
2015 Intake Semester 6 / 2016 Intake Semester 3 Examination

CS5429 DISTRIBUTED COMPUTING

Time allowed: 2 Hours December 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 an open book 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)

An alternative definition for a distributed system is that "*a collection of independent computers providing the view of being a single system*," i.e., it is completely hidden from users that there are multiple computers.

(i) Give an example where this view would become very useful. [3]

(ii) Using a suitable diagram explain how a 3-tier architecture can scale while preserving the single system view. [8]

(iii) Why is it not a good idea to always aim at implementing the highest degree of transparency in distributed systems? Briefly explain. [3]

(iv) Using suitable examples, discuss why a collection of independent computers are unable to provide Consistency, Availability, and Partition Tolerance simultaneously. [6]

(v) "Content Delivery Networks (CDNs) violate the single system view, as they disseminate the content from different points along the request path from client to server"

Do you agree or disagree with this statement? Discuss. [5]

Question 2 (25 marks)

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). Either the same driver or a different driver (box may be exchanged from one truck to another) then delivers the box Rani (Step 4). Finally, the money settlement for the provided service is given to the driver's company (Step 5).

(i) What type of a distributed system architectural style would you recommend for the goods dispatch platform? Briefly justify. [4]

(ii) Which of the following options would you recommend for the communication between the goods dispatch platform and drivers? Justify.

a) Persistent or Transient [3]

b) Asynchronous or Synchronous [3]

(iii) Do you recommend having message queue(s) at the goods dispatch platform? Justify your recommendation. [3]

(iv) Using a suitable diagram illustrate your solution based on recommendations for questions (b) and (c) above. [5]

(v) Do you recommend implementing the goods dispatch platform on a dedicated server(s) in a server room or virtual server(s) on the cloud? Justify. [4]

(vi) Do you recommend having a single server or multiple servers for goods dispatch platform? Justify. [3]

Question 3 (25 marks)

Following is a high-level diagram of a Taxi management system.

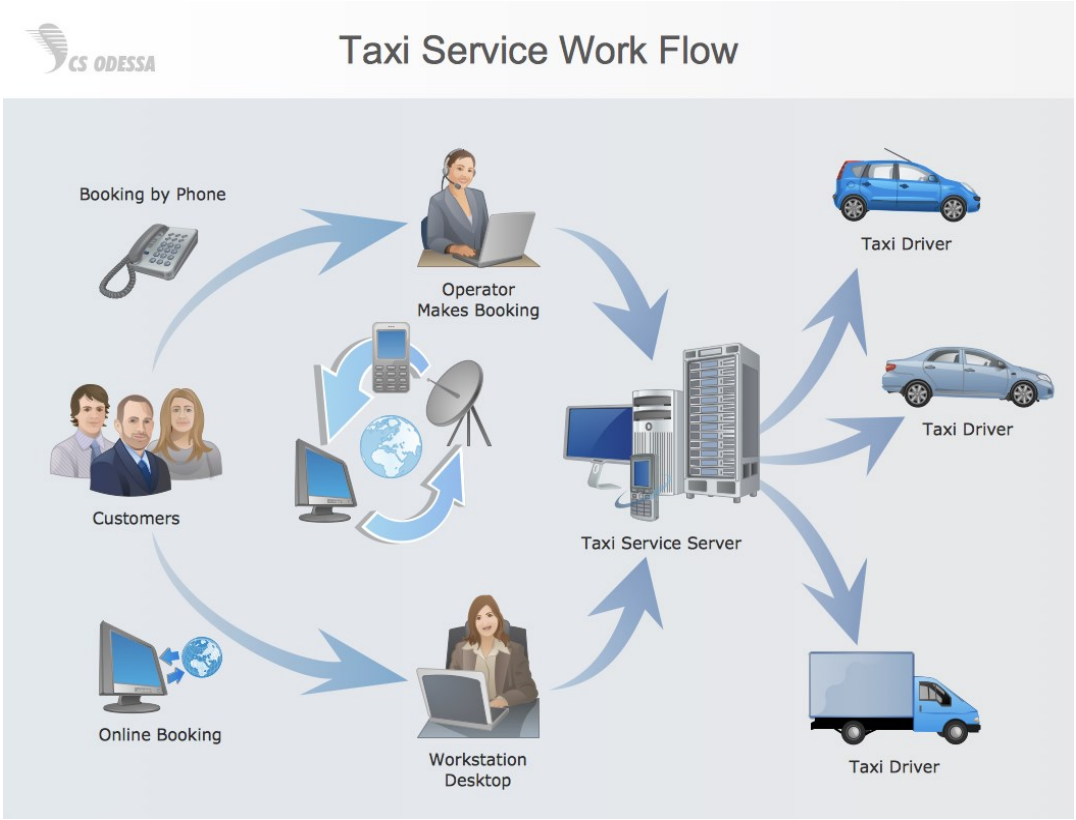

Source: www.conceptdraw.com/samples/business-process-diagrams

Customers who wish to use a taxi may make a booking by calling the operator at the Taxi Service Centre. They may also reserve a taxi via Online Booking, e-mail, or by sending an SMS message with an appropriate format. These messages are received and handled by the Taxi Service Server (TSS). Multiple taxi drivers are registered with the TSS. Customers may use taxies without contacting TSS as well, e.g., if they see a free taxi nearby.

TSS is responsible for assigning a free taxi to a customer, and then notifying both the customer and the taxi driver. Each assignment is based on a several parameters such as (start, destination) of the trip, location of the nearest free taxi, type of vehicle requested (based on AC, non-AC, no of seats, car, van, etc.), and time of the day.

Customers may indicate either an approximate location (e.g., street name and building no) or using GPS coordinates. A GPS device attached to each taxi sends frequent location information to TSS using 3G wireless connection. Taxi drivers also inform the TSS when a taxi becomes free or occupied (i.e., state changed). It is expected that a free taxi be given priority based on first-come-first-serve if they are within 5 Km from a potential customer and other parameters satisfied (e.g., AC, non-AC, no of seats, etc.).

(i) Propose and justify a suitable distributed communication solution for this application.

Hint: you may consider options such as RPC, Web Services, and REST API messages; brokers; and publisher-subscriber. [3]

(ii) Draw a high-level architectural diagram of the proposed solution.

You are expected to clearly indicate different data sources, data consumers, message queues, and key internal and external connectivity. [8]

(iii) How can we enforce first-come-first-serve property for unoccupied taxies? [3]

(iv) Design a protocol for the communication between taxies and TSS. [6]

(v) As a way of reducing the load on TSS servers following suggestion was given by a consultant:

"Directly deliver the taxi request to all nearby taxies (based on GPS coordinates of customer) and let the taxies themselves agree/negotiate on who will cater the customer request"

Discuss the practicality of implementing this solution and its ability to scale. [5]

Question 4 (25 marks)

(i) Assume you have a server that send messages to a client using UDP protocol.

(a) How can you implement at-least-once messaging guarantee for the messages. [4]

(b) Assuming you need to deliver a message to 100,000 clients with at least-once guarantee. Propose and justify a topology to deliver this message most efficiently. [6]

(ii) Department of Inland Revenue wants to rank all the salaried employees from both public and private sector. They have access to monthly salary details of employees. However, it was realized that ranking salaries of a particular month was not a good indication of average income, as monthly salaries of most employees' changes due to no pay leave, over time, bonuses, etc. Hence, it was decided to rank the employees based on average annual salaries over the last 5 years.

(a) Write a pseudocode of a Map-Reduce program that shows how you can calculate the average annual income and rank the employees. The answer should provide all *map* and *reduce* functions, as well as suitable technique to shuffle data between mappers and reducers. [12]

(b) How would you recommend to split the data such that workload of mappers and reducers are balanced? [3]

Question 5 (25 marks)

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

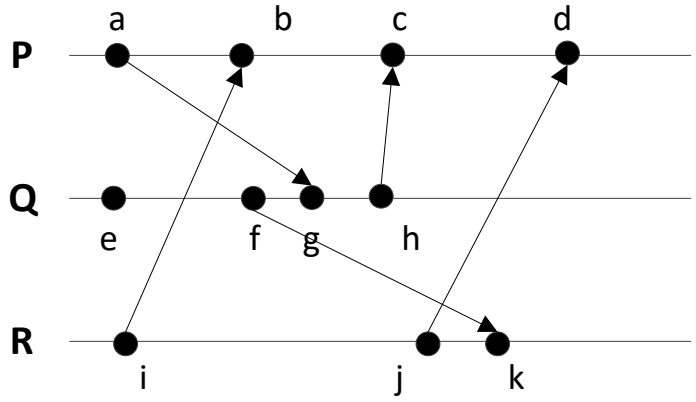

[4]

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 i(b) above? [3]

d) Discuss how one of the event pairs you identified in i(b) above can (or cannot) be ordered using Vector Clocks. [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 3 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) Using an example discuss 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? [4]

(b) Using an example discuss 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? [3]

(c) Using a diagram explain a suitable solution(s) to address issues discussed in ii(a) and ii(b). You may recommend established techniques with justification. [5]

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