UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

B.Sc. Engineering
2010 Intake Semester 8 Examination

CS4262 DISTRIBUTED SYSTEMS

Time allowed: 2 Hours March 2015

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 unauthorised 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)

(i) An alternative definition for a distributed system is that *it is 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. Give an example where this view would become very useful. [4]

(ii) Why is it not always a good idea to aim at implementing the highest degree of transparency possible? Briefly explain. [4]

(iii) Compare and contrast (i.e., similarities and dissimilarities) structured and unstructured overlay networks. [5]

(iv) Amazon Dynamo is a hybrid design between a structured and unstructured overlay network. Discuss how does this hybrid design helps Dynamo to provide fast read access and high availability. [5]

(v) Briefly explain how a key-value store/system can be implemented using Remote Procedure Calls (RPCs). Focus only on storing and retrieving (*key*, *value*) pairs. [7]

Question 2 (25 marks)

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

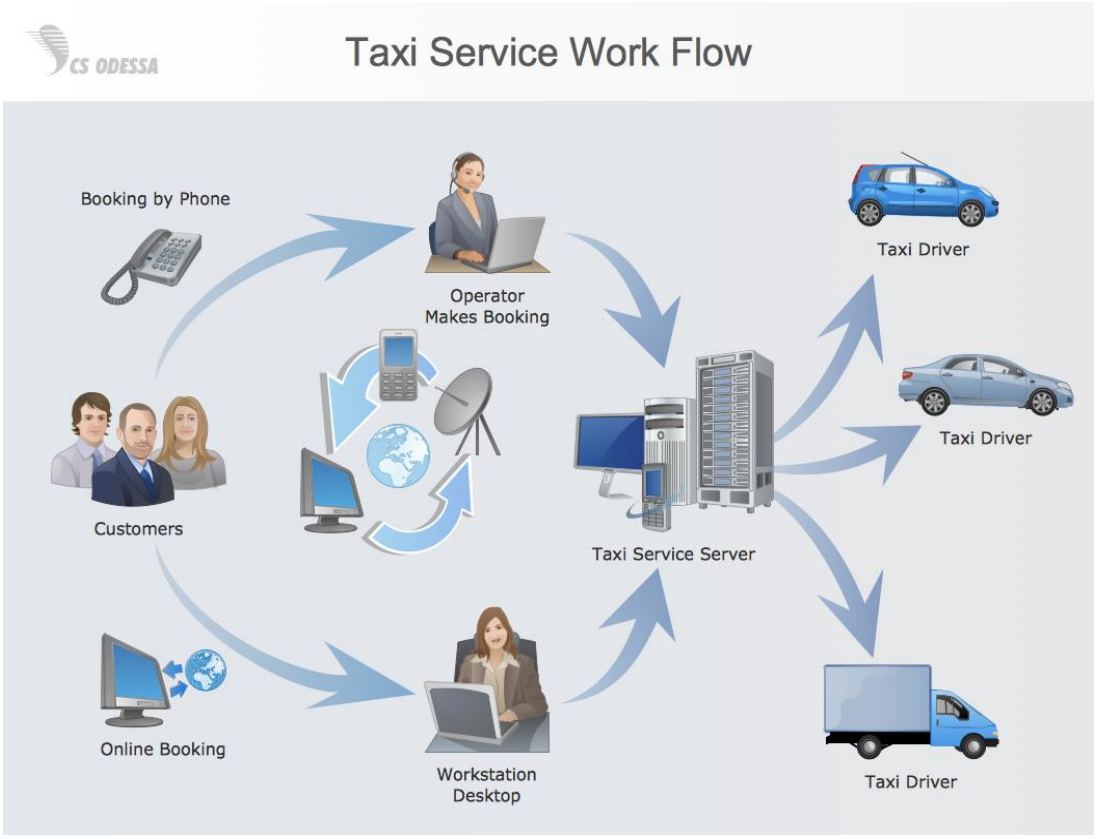

Source: www.conceptdraw.com/samples/business-process-diagrams

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 or by sending an e-mail or 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 the 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) pair of the trip, location of the nearest free taxi, type of vehicle requested (e.g., 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 a building no) or using GPS coordinates. A GPS device attached to each taxi sends frequent location information to the TSS using a 3G wireless data 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 it is within 5 Km from a potential customer and other parameters are satisfied (e.g., AC, non-AC, and no of seats).

a) Do you recommend **Persistent** or **Transient** communication between the customers and TSS? Justify your recommendation. [3]

b) Do you recommend **Asynchronous** or **Synchronous** communication between the customers and TSS? Justify your recommendation. [3]

c) Do you recommend **having a message queue** at the TSS or **not having a message queue**? Justify your recommendation. [3]

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

e) Do you recommend implementing the TSS as a **dedicated server(s) in a server room** or **virtual server(s) on the cloud**? Justify your recommendation. [4]

f) Do you recommend having a **single TSS server** or **multiple TSS servers**? Justify your recommendation. [4]

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

Question 3 (25 marks)

(i) Propose a suitable distributed communication solution for each of the following applications? Justify your suggestion.

(a) A social media aggregation service, that collects a particular user's profile information and posts from multiple social media services such as Twitter, Facebook, Google+, Tumblr, and LinkedIn. This service summarises all interactions of a single user on the same web page. [5]

(b) A music recommendation system that recommends you what songs to listen and what artists to follow based on your music preferences. [5]

(ii) While providing an example, briefly describe any **one** of the following techniques for implementing Content Delivery Networks (CDNs):

- HTTP Redirects based CDN
- Domain Name Service (DNS) based CDN

[5]

(iii) Discuss whether cloud computing is suitable for each of the following applications. Clearly state any assumptions.

(a) To host a website related to the 2015 Cricket World Cup. This website will contain live match updates, player profiles, and related articles. [5]

(b) Sri Lanka's government is trying to introduce a new electronic National Identity Card (namely e-NIC). The e-NIC system is expected to collect personal data and family information. They are also interested in opening up a web service that anyone can use to verify the details of a given NIC number. [5]

Question 4 (25 marks)

(i) Twitter has become the social media solution for rapidly disseminating important news items during social gatherings, disasters, and political events. Twitter users are able to read news that others have never heard about. However, Twitter being owned by a single company creates many issues. For example, a government can demand Twitter to give out all Tweets related to a given event, drop selected Tweets, or they can even block access to the Twitter.com. Privacy concerned users are also worried about all their Tweets being accumulated by a single entity.

A team of volunteers is starting a project to address these issues in Twitter by developing a distributed system. They are envisioning a system where there is no single controlling entity and users will contribute to keep the system running, e.g., by contributing processing and storage resources as in peer-to-peer file sharing systems. The volunteers have decided to support both public and user-to-user Tweets.

a) What type of an overlay network would you recommend for this solution? Justify your recommendation. [3]

b) Propose a suitable naming scheme to uniquely identify each Tweet. [3]

c) What type of a communication mechanism (network sockets, RPC, RMI, web services, etc.) would you recommend for this solution? Justify your recommendation. [3]

d) Using a suitable diagram explain your overall solution to disseminate the Tweets among users/nodes. [6]

e) Twitter has recently started attaching photos to the Tweets. Discuss how your solution can be extended to support disseminating photos attached to the Tweets. [3]

f) Explain “Replication Transparency” in the context of your proposed solution. [3]

g) Discuss to what extent your solution meets the desired security and privacy goals. [4]

Question 5 (25 marks)

(i) Why clock drift and skew are problematic in distributed systems? Briefly explain. [4]

(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.

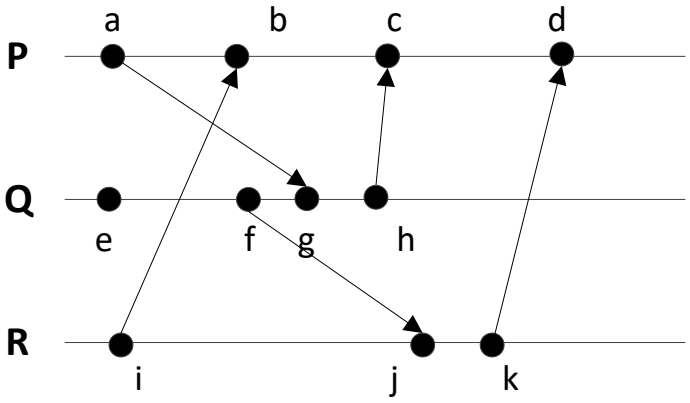

[5]

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? [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.

(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) Using a diagram explain a suitable solution to address these issues. You may recommend established techniques with justification. [6]

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