

Index No:							
-----------	--	--	--	--	--	--	--



UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

B.Sc. Engineering

2012 Intake Semester 8 Examination

CS4262 DISTRIBUTED SYSTEMS

Time allowed: 2 Hours

March 2017

ADDITIONAL MATERIAL: *None*

INSTRUCTIONS TO CANDIDATES:

1. This paper consists of **12** pages.
2. Answer **ALL** questions from **Section A** and **any three (3)** questions from **Section B**.
3. Answer the questions on the paper itself. Answer within the given space.
4. For MCQ and True/False questions, select the most appropriate answer. No penalty for wrong answers.
5. The maximum attainable mark for each question is given in brackets.
6. This examination accounts for 50% of the module assessment.
7. This is a closed book examination.

NB: It is an offence to be in possession of unauthorised material during the examination.

8. Only calculators approved by the Faculty of Engineering are permitted.
9. Assume reasonable values for any data not given in or with the examination paper. Clearly state such assumptions made on the script.
10. 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.
11. This paper should be answered only in English.

Q1	Q2	Q3	Q4	Q5	Q6	Q7	Total

Section A**Question 1****[2 × 10 = 20 marks]**

Circle the most appropriate answer.

- (i) Which of the following statements is **not true** about distributed systems?
- a. Attempts to provide a Single System View to users
 - b. Built on a homogeneous infrastructure
 - c. Relatively easy to expand or scale
 - d. Use message passing for communication and coordination
- (ii) Migration transparency
- a. Allows access without knowledge of location
 - b. Enables multiple instances of resources
 - c. Enables the movement of resources
 - d. Enables the concealment of faults
- (iii) Which of the following is **not** a basic scaling technique for distributed systems?
- a. Distribution
 - b. Heterogeneity
 - c. Hiding latency
 - d. Replication
- (iv) Which of the following are true about P2P networks?
- (p) Unstructured networks are relatively easier to setup.
 - (q) Unstructured networks are more suitable to handle mutable objects.
 - (r) Structured networks provide bounded performance.
- a. *p* and *q* only
 - b. *p* and *r* only
 - c. *q* and *r* only
 - d. All three
- (v) Which of the following statements is true about P2P streaming?
- a. It is relatively easier to setup a tree-push network.
 - b. Tree-push approach has higher resilience than mesh-pull approach.
 - c. Mesh-pull approach has the lowest latency.
 - d. Mesh-pull approach uses a bit map to identify video frames on peers.
- (vi) Which of the following are true about directory services?
- (p) Map names to attributes.
 - (q) Support only service discovery.
 - (r) Used to identify RPC and Web services.
- a. *p* and *q* only
 - b. *p* and *r* only
 - c. *q* and *r* only
 - d. All three

- (vii) What feature **does not** belong to a private cloud?
 - a. Dial-home support
 - b. Metered billing
 - c. Self-service portal
 - d. Rapid elasticity

- (viii) Physical time in a distributed system does not make much sense because
 - a. Clock skew is almost constant due to small differences in drift rate.
 - b. Asynchronous models are sufficient in distributed systems.
 - c. There is an upper bound on message transmission delays.
 - d. We cannot synchronize clocks on multiple nodes.

- (ix) Which of the following property of transactions ensures that concurrent transactions do not interfere with each other?
 - a. Atomic
 - b. Consistent
 - c. Durable
 - d. Isolated

- (x) Which of the following are true about transactions?
 - (p) Dirty Read problem can be solved using Serial Equivalence.
 - (q) Lost Update problem can be solved using Serial Equivalence.
 - (r) Totally ordered multicast addresses concurrency problem.
 - c. p and q only
 - d. p and r only
 - c. q and r only
 - d. All three

Question 2

[1 × 10 = 10 marks]

Tick True or False.

		True	False
(i)	Web services are an example of Object-Based Architectures.		
(ii)	Flooding in unstructured networks guarantees content discovery.		
(iii)	In RMI, state of remote objects is distributed.		
(iv)	Web caching is an example of achieving best effort consistency.		
(v)	Split-brain syndrome can be fixed by having three instances of a given service.		
(vi)	Response-based, transient, synchronous communication is weaker than receipt-based, transient, synchronous communication.		
(vii)	Message queues could increase response time.		
(viii)	Row versioning provides a centralized mechanism to provide isolation among transactions.		

		True	False
(ix)	Separation of file system structure (i.e., name space) and the contents of files, enables Google file system to scale to petabytes of data while having megabytes of metadata at the master node.		
(x)	If a distributed system has 2 nodes each with availability of 99.9, overall availability of the system is 99.9^2 .		

Question 3**[1 × 10 = 10 marks]**

Fill in the blanks using one of the following keywords:

Asynchronous, Bootstrapping, Broker, Data manager, Distributed, FIFO, Flat, Horizontal, Hybrid, Hypervisor, Multicast, Overlay, Rarest-first, Rapid-development, Scalability, Scheduler, Transient, Underlay, Vertical, Wait free, Xen

- (i) In _____ scaling, we gain better performance by using a machine with more resources (e.g., CPU and memory).
- (ii) _____ is an application layer network, built on top of a physical network.
- (iii) _____ is the process used by new node in a distributed system to identify other nodes in the system.
- (iv) _____ dissemination policy in BitTorrent enables rapid dissemination of contents.
- (v) In _____ communication, sender continues immediately after it has submitted its message for transmission.
- (vi) A message _____ decouples publishers from subscribers.
- (vii) _____ cloud computing environments use a mix of on-premises, private cloud and third-party, public cloud services.
- (viii) _____ transaction is logically a flat, indivisible transaction that operates on distributed data.
- (ix) One key advantage of the Platform as a Service (PaaS) cloud model is _____.
- (x) _____ is responsible for proper control of concurrency in transactions.

Section B

Answer any three (3) questions.

Question 4

[20 marks]

- (i) Describe 2 benefits of replication in a video distribution service, e.g., Youtube. [4 marks]

- (ii) What kinds of consistency problems can the use of replication cause? Discuss. [4 marks]

- (iii) Propose, with justification, a reasonable consistency model for the video distribution service. [4 marks]

- (iv) Which active Content Delivery Networks (CDN) option is most suitable for for the video distribution service? Justify. [4 marks]

- (v) Suppose that you have access to only transient, asynchronous communication primitives. What changes are required to provide primitives for transient, synchronous communication? [4 marks]

Question 5**[20 marks]**

- (i) Which property(ies) among Consistency, Availability, and Partition Tolerance is affected in each of the following cases: [1 × 4 = 4 marks]

a)	Split brain syndrome	
b)	Amazon Dynamo	
c)	An ATM that refuses to give money as it cannot talk to core backing system	
d)	Accessing Google Drive without connecting to Internet	

- (ii) When a transaction is aborted, we have said that the world is restored to its previous state, as though the transaction had never happened. Give an example where resetting the world is impossible when a transaction is aborted. [4 marks]


- (iii) Consider a multiplayer cards game in which different players take turns in the order in which they joined the game. Design such a game on a server-less architecture (i.e., without any central control). Clearly state any assumptions. [9 marks]

- (iv) If few players got disconnected, discuss the ability of your solution to continue with the connected players. [3 marks]

Question 7

[20 marks]

- (i) Label the following diagram with Lamport's Timestamps. [4 marks]



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

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.

- (iii) 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 marks]

- (iv) How would you recommend to split the data such that workload of mappers and reducers are balanced? [2 marks]

Question 7

[20 marks]

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 is to be developed. A person’s activity and eating habits are to be tracked via smart watches, smart phones, tablets, wireless weight scales, desktops, etc. This service is also to be coupled with a health care network of doctors, hospitals, and Ambulance services. Two high-level feature diagrams are illustrated below.



Source: adamelements.com

(i) Tick True or False.

[1 × 6 = 6 marks]

		True	False
a)	Event-based architectural style would be suitable for this platform.		
b)	Transient communication is suitable for communication between wearable/handheld devices and the backend.		
c)	Asynchronous communication is suitable for communication between wearable/handheld devices and the backend.		
d)	When the bandwidth between wearable/handheld devices and the backend is limited, it is acceptable to randomly drop messages/packets.		
e)	Dynamic load balancing is more appropriate at the backend.		
f)	Publisher-subscriber mode of message dissemination is suitable to provide health news.		

(ii) What type of a cloud implementation would you recommend to implement this solution? Justify considering options such as IaaS, PaaS, SaaS, and containers. [4 marks]

(iii) What do you recommend, RPC, RMI, REST API, or Web services, to integrate these services? Justify. [4 marks]

- (iv) Draw a high-level architectural diagram illustrating the key components and their connectivity. [6 marks]



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