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

**UNIVERSITY OF MORATUWA**

Faculty of Engineering

Department of Computer Science & Engineering

B.Sc. Engineering
2014 Intake Semester 8 Examination

**CS4262 Distributed Systems**

Time allowed: 2 Hours December 2018

---

**ADDITIONAL MATERIAL:** *None*

## INSTRUCTIONS TO CANDIDATES:

1. This paper contains **4** questions on **12** pages.
2. This examination accounts for 50% of the module assessment.
3. This is a closed book examination.
4. Answer **ALL** questions.
5. Answer the questions on the paper itself and within the given space.
6. For MCQ and True/False questions, select the most appropriate answer. No penalty for wrong answers.
7. The maximum attainable mark for each question is given in brackets.
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. If you have any doubt as to the interpretation of the wording of a question, make your own decision, but clearly state it on the script.
11. Electronic/Communication devices are not permitted. Only equipment allowed is a calculator approved and labelled by the Faculty of Engineering.
12. This paper should be answered only in English.

| Q1 | Q2 | Q3 | Q4 | Total |
|---|---|---|---|---|
| | | | | |

**Question 1** **[25 marks]**

Circle the most appropriate answer. [2 × 10 marks]

(i) Which of the following statements does **not** reflect the characteristics of a distributed system?

a. Maintains global state despite of node churn and failures

b. Presents a single system view to users while hiding the existence of multiple nodes

c. Handles varying loads via horizontally scaled stateless components

d. Tradeoff consistency to maintain availability during network partitioning

(ii) Which of the following is a **not** a characteristic of Cloud Computing?

a. Centralized

b. Rapid elasticity

c. Service orientation

d. Rely on virtualization

(iii) Which of the following statements are true?

(p) Network File System is an example of Access Transparency.
(q) Google Drive is an example of Location Transparency.
(r) Lankadeepa.lk is an example of Replication Transparency.

a. *p* and *q* only

b. *p* and *r* only

c. *q* and *r* only

d. All three

(iv) In a modern multi-tiered web application

a. Each tier can be scaled independent of others

b. Overall availability is determined by tier with the highest numbers of nodes

c. Rely only on host-level load balancing

d. Typically runs active-active databases

(v) Which of the following statement is true about P2P networks?

a. Unstructured networks provide guaranteed content discovery with predictable bounds

b. Unstructured networks are desirable in small-scale or highly dynamic systems

c. Structured networks could tolerate higher level of peer churn and failure

d. Structured networks are desirable in indexing mutable objects with multiple complex attributes

(vi) Which of the following is true about distribute mutual exclusion algorithms

a. Centralized algorithm guarantees first come first serve property

b. Decentralized algorithm works well under high contention

c. Token-based algorithm does not scale as worst case waiting time is $O(n)$

d. Distributed algorithm requires physical clocks of nodes to be in sync

(vii) Which of the following is true about distribute mutual exclusion algorithms

a. Bully algorithm can scale to large networks with high churn and failures
b. Ring algorithm can scale to large networks with high churn and failures
c. Require nodes identifiers to be globally unique
d. Require peers to know all other peers in the system

(viii) To support failure atomicity

a. Needs synchronize operations
b. Partial results need to be visible
c. Operations must be indivisible
d. Objects must be recoverable

(ix) Which of the following statements is true about blockchain?

a. Consensus based on an elected leader provides weak consistency
b. Consensus based on an elected leader provides consistency even in a partitioned network
c. Proof of Work (PoW) based consensus provides eventual consistency
d. Average time to solve the puzzle in Proof of Work (PoW) increase with number of blockchain nodes

(x) Which of the following is true about transactions?

a. Dirty Read problem can be solved using Serial Equivalence
b. Impact of Dirty Read problem can be recovered using Write-Ahead Log
c. Lost Update problem can be solved using Serial Equivalence
d. Nested transactions work on distributed data

(xi) Fill in the blanks using one of the following keywords: [1 × 5 marks]

```
Active | Batch | Byzantine | f - 1 | f + 1 | Nakamoto |
Passive | Request-level | Stream | Thread-level
```

a. High throughput in Cloud Computing is achieved via _____ parallelism.

b. _____ caching is still effective in the presence of HTTPS.

c. _____ processing operates on data in motion.

d. Agreement/Consensus in the presence of misbehaving nodes need _____ fault tolerance.

e. To reach consensus with *f* crash failures requires at least _____ rounds of message exchanges.

**Question 2** **[25 marks]**

(i) Using an example, briefly explain how scaling transparency could be achieved better via stateless components. [5 marks]

Consider the following design (Figure 1) of a web-based application. Availability of each component and service is given in the figure. To enhance the throughput and reliability of the web-based application, it is proposed to run 2 instances of both the Web Role and Worker Role.

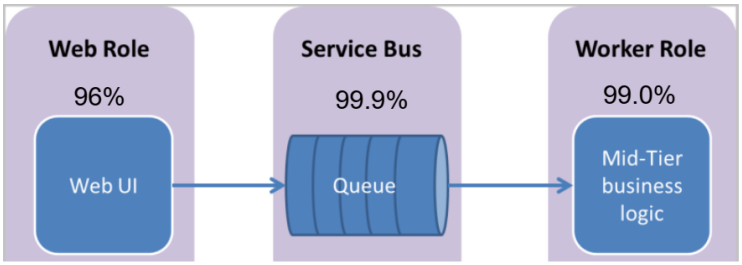


Figure 1 – 3-tier web application.

(ii) What component or service is the most difficult to scale? Justify. [3 marks]

(iii) What component or service is most likely to fail? Justify. [2 marks]

(iv) Using a rough calculation, explain how the overall availability of the application would improve if 2 instances of Web Role and Worker Role are used instead of one. [5 marks]

(v) While retaining the Service Bus how could we split the workload among the 2 Worker Roles? Discuss. [5 marks]

(vi) When 2 nodes run the Worker Role, should we increase the queue length of Service Bus compared to the case of having only a single instance of Worker role? Discuss. [5 marks]

**Question 3** **[25 marks]**

The Things Network (TTN) attempts to build an overlay network using Internet of Things (IoT) devices that are around us. Like the Internet (which is a collection of interconnected networks), TTN plans to build a network of IoT networks. Using TTN we should be able to communicate with other things without the reliance of infrastructure based-networks such as 3G/4G and Ethernet by relaying on P2P routing.

Figure 2 shows the high-level architecture of TTN. LoRaWAN is a very-long-range and low-power wireless technology that enables communication among IoT devices. Join Server safely stores LoRaWAN keys and issues session keys to the Network Server and Application Server. Identity server provides device authentication.

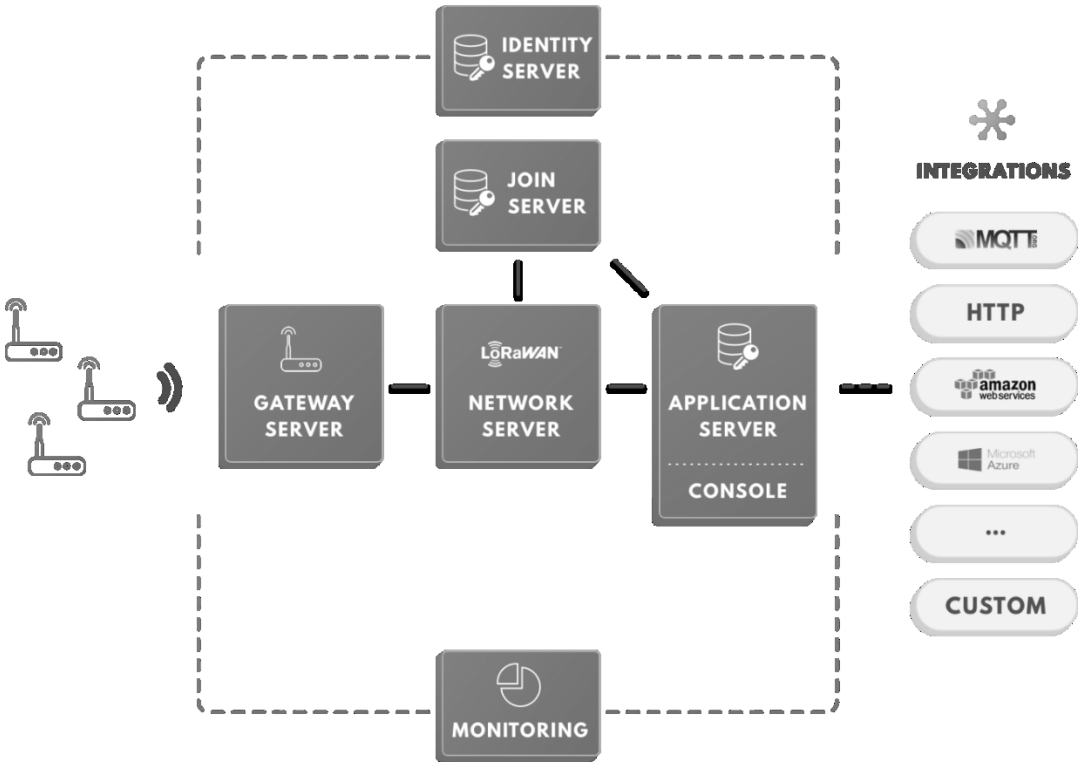


Figure 2 – High-level architecture of Things Network. Source: www.thethingsnetwork.org

(i) Among Consistency, Availability, and Partition Tolerance (CAP) which property would you propose to compromise in TTN? Discuss. [3 marks]

(ii) What overlay topology would you propose to interconnect IoT nodes? Justify. [4 marks]

(iii) Would you recommend, a publisher/subscriber service or an Enterprise Service Bus (ESB) for the integration with external data consumption services such as AWS and Google IoT? Discuss. [4 marks]

(iv) Should we time stamp the messages at the IoT device, Gateway Server, or Network Server? Discuss. [7 marks]

(v) Suppose IoT devices/nodes could receive micro payments for contributing to the TTN. What particular challenges need to be overcome to support distributed transactions in this case? Discuss. [4 marks]

(vi) Tick True or False. [1 × 3 marks]

| | | True | False |
|---|---|---|---|
| (a) | Event-based architectural style would be suitable for this platform. | | |
| (b) | Join Server is stateful. | | |
| (c) | Transient communication is suitable for the communication between IoT devices and TTN servers. | | |

**Question 4** **[25 marks]**

(i) "*Content Delivery Networks (CDNs) enhance performance while compromising consistency*." Do you agree or disagree with this statement? Discuss. [4 marks]

(ii) "*Virtualization is not the same as Cloud Computing.*" Do you agree or disagree with this statement? Discuss. [4 marks]

Suppose Twitter is interested in generating a Word Cloud of hashtags (see Figure 3) that appeared in the last 15 minutes. *Word Cloud* is an image composed of words used in a particular text, in which the size of each word indicates its frequency. Thus, we need to find the frequency of hashtags that appeared in the last 15 min. Suppose Twitter plans to plot only the 100 most popular hashtags.


Figure 3 – Word Cloud. Source: arnesund.com

(iii) Write a pseudocode of a Map-Reduce program that shows how you can calculate the ranked frequency table of hashtags. The answer should provide all map and reduce functions, as well as technique to shuffle data between mappers and reducers. [11 marks]

(iv) Suppose Twitter wants to update the Word Cloud every 30 seconds based on the Tweets appeared in the last 5 min. Draw an Apache Storm style topology (DAG of spouts and bolts) to calculate the rolling top hashtags. Label the diagram. [6 marks]

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