NAMIBIA UNIVERSITY OF SCIENCE AND TECHNOLOGY

FACULTY OF COMPUTING AND INFORMATICS

DEPARTMENT OF COMPUTER SCIENCE

QUALIFICATION: BACHELOR OF COMPUTER SCIENCE	
QUALIFCATION CODE: 07BACS	LEVEL: 6
COURSE: DISTRIBUTED SYSTEMS PROGRAMMING	COURSE CODE: DSP620S
DATE: November 2019	SESSION: 1
DURATION: 3 Hours	MARKS: 100

FIRST OPPORTUNITY EXAMINATION QUESTION PAPER	
EXAMINER:	Prof. José G. Quenum
MODERATOR:	Prof. Dharm Singh Jat

This paper consists of 3 pages
(excluding this front page)

INSTRUCTIONS

1. This paper contains 6 questions.
2. Answer all questions on the exam paper.
3. Marks/scores are provided at the right end of each question
4. Do not use or bring into the examination venue books, programmable calculators, mobile devices and other materiaks that may provide you with unfair advantage. Should you be in possession of one right now, draw the attention of the examiner officer or the invigilator.
5. NUST examination rules and regulations apply.

PERMISSIBLE MATERIALS

None

Question 1 ..... [15 points]

Consider a `TaskRabbit` application, where a client submits a job for execution. The client submits the job to an agent, who then delegates it to a worker. The overall protocol goes as follows: the client sends the description of the task to an agent. The latter then forwards the description to a worker. When the task is completed sends the notification back to the agent, who then forwards it to the client. The three (3) participants use two (2) socket-based protocols. The client, the agent and the worker all use a TCP-based socket.

Extend the TCP-based socket diagram presented in class to represent the overall protocol.

Question 2 ..... [30 points]

(a) Consider a `kafka` cluster containing five (05) brokers, CL_1, CL_2, CL_3, CL_4 and CL_5. Each topic contains four (04) partitions with a replication factor of three (03), i.e. each partition is replicated twice (on different brokers). Using a diagram representing the cluster illustrate how a producer submits messages to the cluster, and a consumer group consumes such messages. You will be explicit about how the partitions are handled. [20]

(b) Consider the diagram in Figure 1. It depicts the infrastructure for a publish-subscribe communication mechanism with one publisher, eight (8) subscribers and a network of brokers. [10]

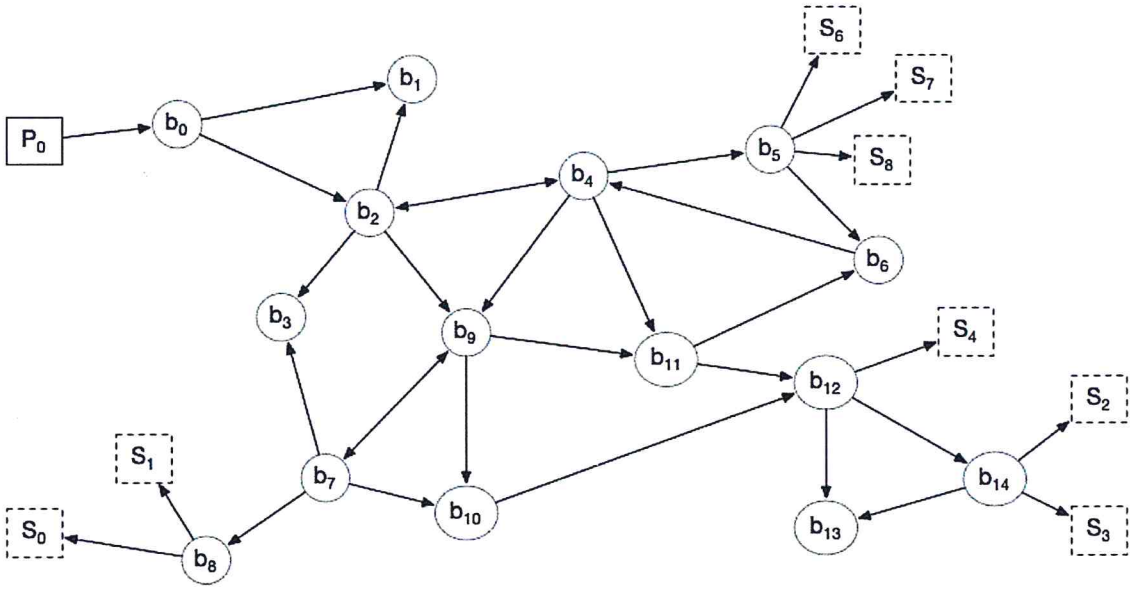

Figure 1: A Publish-Subscribe Infrastructure

Subscribers S_0 and S_8 subscribed for a topic named DSPCourse1. Table 1 is the portion of the routing table availed to individual brokers in the network. Using the filter routing technique show the itinerary of two messages m_1 (topic: DSPCourse) and m_2 (topic: DSPCourse1) sent by publisher P_0 to all potential subscribers who should receive the message.

Table 1: Routing Table

Broker	Destination	Topic
b_0	b_2	ANY
b_2	b_4	DSPCourse, ANY
	b_9	DSPCourse
b_9	b_7	DSPCourse
	b_{11}	ANY
b_7	b_8	DSPCourse
	b_9	ANY
b_4	b_5	DSPCourse
	b_2	ANY
	b_{11}	ANY
b_{11}	b_{12}	ANY
b_{12}	b_{13}	DSPCourse1

Question 3 ..... [15 points]

Two processes S_i and S_j are involved in remote invocation with S_i acting as the caller, while S_j acts as the callee.

(a) Highlight the underlying request-reply protocol using its primitives. [6]

(b) The request issued by S_i was not replied to on time. Explain the steps taken by both processes using the at-least-once semantics. [9]

Question 4 ..... [20 points]

Building on the abstract architecture defined for distributed file systems in class, we wish to design a distributed storage where we can store objects represented by a key and a (binary) value. The additional requirements for the system are as follows:

- transparency in access and location;
- concurrent operations, including updates;
- replication, fault-tolerance and concurrency.

Using a diagram, propose an architecture for our distributed storage. You will explain how each requirement is fulfilled.

Question 5 ..... [10 points]

Consider the table below depicting several processes in a distributed system and the resource they use in an election algorithm. The table shows, for each process, the process identifier, and the quantity of the resource the election is based on.

Process Id.	Resource	Process Id.	Resource
P_0	17	P_1	24
P_2	1	P_3	28
P_4	15	P_5	9
P_6	4	P_7	3

The current leader, P_3 has just crashed. The remaining processes start an election using the bully algorithm after detecting the crash on P_3. However, process P_1 crashes right after receiving to the first message. Describe step-by-step the execution of the leader election and the winner.

Question 6 ..... [10 points]

Consider four (04) processes (P_1, P_2, P_3 and P_4) as part of a distributed system. The following events (listed in order) have occurred at each process:

P_1: e_0, e_1, e_2, e_3, e_4 and e_5

P_2: i_0, i_1, i_2, i_3 and i_4

P_3: j_0, j_1, j_2 and j_3

P_4: ℓ_0, ℓ_1, ℓ_2 and ℓ_3

In addition, we have the following observations:

- Event i_1 resulted from a message exchange between P_1 and P_2 after event e_2;
- Event e_4 resulted from a message exchange between P_2 and P_1 after event i_2;
- Event j_0 resulted from a message exchange between P_1 and P_3 after event e_1;
- Event i_4 resulted from a message exchange between P_3 and P_2 after event j_1;
- Event e_3 resulted from a message exchange between P_4 and P_1 after event ℓ_0;
- Event ℓ_1 resulted from a message exchange between P_3 and P_4 after event j_2;
- Event e_5 resulted from a message exchange between P_4 and P_1 after event ℓ_2;
- Event ℓ_3 resulted from a message exchange between P_3 and P_4 after event j_3.

Using a diagram represent the vector clocks corresponding to the logical clock for each event in the system.