

Family Name:..... First Name:.....

Student ID:..... Signature

NWEN 304 : Mid-term Test I

18 August 2023 ** WITH SOLUTIONS **

Instructions

- Time allowed: **50 minutes**
- Attempt **all** the questions. There are 50 marks in total.
- **In-person:** Write your answers in this test paper and hand in all sheets.
Remote: Type your answers in the template file and submit to **Remote-Test-I** on the NWEN 304 submission system.
- If you think a question is unclear, ask for clarification.
- This test contributes 25% of your final grade.
- You may use dictionaries and calculators.
- You may write notes and working on this paper, but make sure your answers are clear.

Sections

Marks

1. True/False	[10]	<input type="text"/>
2. Multiple Choice	[10]	<input type="text"/>
3. Multiple, Multiple Choice	[10]	<input type="text"/>
4. Short Answer	[20]	<input type="text"/>
	TOTAL:	<input type="text"/>

Question 1. State True or False. Write your answer in the box provided.

[10 marks]

- (a) [1 mark] Client-Server architecture does not allow direct communication between clients.

True or T

- (b) [1 mark] The components of a distributed system have the ability to communicate with each other through a shared physical memory.

False or F

- (c) [1 mark] The concept of Watchers in Zookeeper is based on the notion of client polling.

False or F

- (d) [1 mark] An advantage of employing asynchronous communication via message brokers is that it enables services to exchange information even in instances where some of them are temporarily unavailable.

True or T

- (e) [1 mark] For each new incoming request, Node.js spawns a new thread.

False or F

- (f) [1 mark] Including multiple external JavaScript files into an HTML document could lead to the initiation of multiple HTTP requests to the server for fetching these files.

True or T

- (g) [1 mark] One of the reasons for refraining from incorporating business logic in the Presentation tier of a 3-tier architecture is that the code executes in the client browser, which is both accessible and visible to the user.

True or T

- (h) [1 mark] In ZooKeeper, by default all zNodes are Ephemeral unless otherwise stated.

False or F

- (i) [1 mark] Distributed Hash Tables are used to support look-up operation in unstructured P2P networks.

False or F

- (j) [1 mark] One of the reasons for the popularity of Node.js over Java in the development of data-intensive applications is Node.js's utilization of blocking I/O, in contrast to Java's use of non-blocking I/O.

False or F

Question 2. Multiple choice questions. Each question has a single correct option. [10]
Write your answer in the box provided.

(a) [1 mark] What does the following JavaScript code log?

```
console.log(typeof [ ]);
```

- (i) Array
- (ii) Object
- (iii) Undefined
- (iv) Null

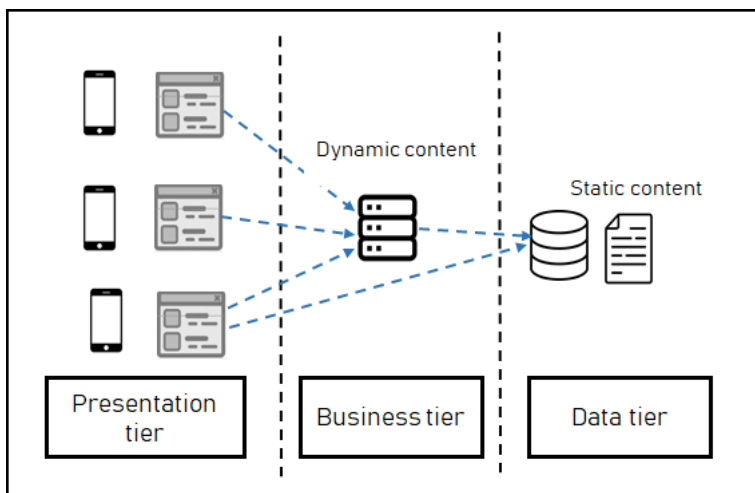
ii

(b) [1 mark] Which one of the following operations is **NOT** idempotent?

- (i) Deleting a record with a particular ID from the database.
- (ii) Resetting a user's balance to 0.
- (iii) Evaluating the absolute value of a number.
- (iv) Removing the last record from a queue.

iv

(c) [1 mark] Considering the architecture of the client-server application provided below, choose the appropriate statement from the following options.



- (i) This is a well designed 3-tier architecture diagram.
- (ii) This architecture is problematic as it contains both files and a database in the Database tier.
- (iii) This architecture is problematic as the Presentation tier bypasses the Business tier and sends requests directly to the Data tier.
- (iv) This architecture is problematic as both (ii) and (iii) are correct.

iii

(d) [1 mark] Why is a 3-tier Architecture generally referred to as a Monolithic Architecture?

- (i) Because all the business logic is concentrated in one deployable unit.
- (ii) Because the entire frontend code runs on a single client device.
- (iii) Because we can have only one database in the Data tier.
- (iv) Because we can have only one codebase, and we can't break it into multiple modules.

i

(e) [1 mark] In which of the following two protocols is data read as a byte-stream?

- (i) UDP
- (ii) TCP

ii

(f) [1 mark] What will be the output of the following JavaScript statement: `45 % -2`?

- (i) NaN
- (ii) -1
- (iii) 1
- (iv) Error

iii

(g) [1 mark] Which one of the following P2P architectures is best suited for supporting distributed cache?

- (i) Structured P2P
- (ii) Flat unstructured P2P
- (iii) Loosely structured P2P
- (iv) Two level unstructured P2P

i

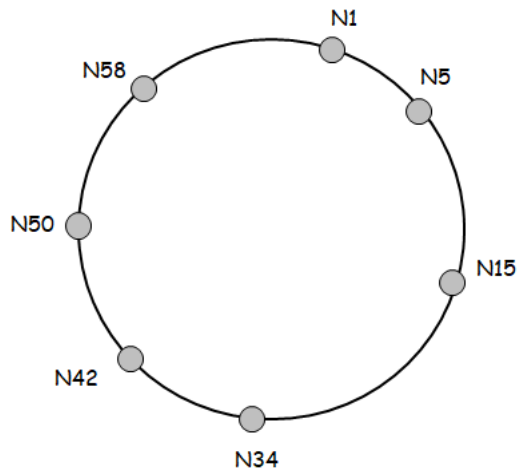
(h) [1 mark] What is the correct JavaScript syntax to change the content of the HTML element below?

```
<p id="demo">This is a demonstration.</p>
```

- (i) `document.getElementByName("p").innerHTML = "Hello World!";`
- (ii) `document.getElementById("demo").innerHTML = "Hello World!";`
- (iii) `document.getElement("p").innerHTML = "Hello World!";`
- (iv) `#demo.innerHTML = "Hello World!";`

ii

(i) [1 mark] Consider the following representation of a consistent hash ring with a hash space of size 2^6 . Where should the key k_{60} be assigned?



- (i) N58
- (ii) N1
- (iii) k_{60} cannot be assigned as the node with hash key 60 is not active.
- (iv) k_{60} cannot be assigned as no node with hash key 0 is active.

ii

(j) [1 mark] Fill in the blank. In a P2P Chord network, the finger table maintained by a node holds a set of _____ entries for its successor node IDs.

- (i) $\log(N)$, where N is the number of active nodes in the network.
- (ii) $\log(N)$, where N is the number of keys to be stored.
- (iii) $\log(N)$, where N is the size of the hash space.
- (iv) $\log(N)$, where N is the number of inactive nodes in the network.

iii

Question 3. Multiple, multichoice questions. Each question has multiple correct options.[10]
Write your answer in the box provided.

- (a) [2 marks] Which of the following statements are true about Apache Zookeeper?
- (i) Zookeeper cluster follows a leader follower model.
 - (ii) To avoid single point of failure, Zookeeper is typically run as a cluster of n nodes where n is any integer greater than 1.
 - (iii) Ephemeral Znodes can serve as a means to store application-wide information that could be utilized by all active nodes within the computing cluster.
 - (iv) Zookeeper clients establish a connection to Zookeeper using the TCP protocol.

i,iv

- (b) [2 marks] Which of the following statements are true for consistent hashing?
- (i) When a new node n joins, it necessitates the transfer of certain keys from node n 's predecessor to node n .
 - (ii) When a new node n joins, it necessitates the transfer of certain keys from node n 's successor to node n .
 - (iii) When a node n leaves, it necessitates the transfer of certain keys from node n to n 's predecessor.
 - (iv) When a node n leaves, it necessitates the transfer of certain keys from node n to n 's successor.

ii,iv

- (c) [2 marks] Which of the following statements are incorrect in relation to TCP and UDP protocols?
- (i) TCP provides throughput guarantees.
 - (ii) Both TCP and UDP provides throughput guarantees.
 - (iii) None of TCP and UDP provides throughput guarantees.
 - (iv) TCP guarantees delivery of data.

i,ii

- (d) [2 marks] Which of the following statements are true for API design?
- (i) RPC based APIs are tightly coupled with the service implementation.
 - (ii) An advantage of Resource API over Message API is that Resource APIs use a small number of standardized methods to support various operations while Message APIs require distinct messages for each combination of domain entity and the associated operation.
 - (iii) Asynchronous communication involves sending a request and blocking while waiting for the response.
 - (iv) Fetch API is a synchronous API.

i,ii

(e) [2 marks] Which of the following statements are true for Node.js as a server side JavaScript run-time environment?

- (i) Node.js uses a non-blocking I/O model.
- (ii) Node.js is equally good for data intensive and compute intensive applications.
- (iii) Node.js HTTP interface does not buffer entire requests or responses.
- (iv) Node.js HTTP interface buffers entire request and responses.

i,iii

Question 4. Short answer questions**[20 marks]**

- (a) **[4 marks]** Explain the constituent architectural elements of an application's architecture.

Components – units that provides a transformation of data via its interfaces
(run as separate runtime units)

Connectors – mechanism that allows communication, coordination and co-operation among components

Data – element of information that is transferred/received by a component via a connector

Interfaces – Connection points on components and connectors

(b) [4 marks] Suppose you are implementing a request-response service. Explain network error handling functionalities that you would implement (at the client and/or the server) to support the following operations:

- (i) PUT request: The PUT request was unsuccessful in reaching the server due to a transient network issue.
- (ii) POST request: The server effectively handled the POST request, the response to the POST couldn't reach the client due to a transient network error.

1. PUT request: wait for certain period of time, if the response is not received resend the request

2. POST request: Since no response has been received by the client, the client may infer two possible scenarios:

a) the POST request may have failed and the server did not service the request or

b) the POST request was successful but the acknowledgment was lost

In either case, the client wouldn't be able to determine which of the two scenarios led to the lack of acknowledgment reception and hence has no choice but to resend the request after a timeout. A mechanism needs to be established on the server side to guarantee that the POST operation isn't duplicated if it succeeded initially. The implementation of the POST should be designed to be idempotent, which could be achieved through the utilization of unique IDs.

(c) [6 marks] Explain the following in relation to cluster coordination.

(i) [2 marks] The Herd effect.

(ii) [4 marks] How is the Herd effect managed in the implementation of leader election in a compute cluster using Zookeeper as the cluster coordinator?

The herd effect refers to releasing a "herd" of requests from multiple machines when in fact only a single or a small number of machines can proceed.

For fault tolerance, it is important to watch for failures of the leader, so that a new client arises as the new leader in the case the current leader fails.

A trivial solution is to have all application processes watching upon the current smallest znode, and checking if they are the new leader when the smallest znode goes away. But this causes a herd effect: upon failure of the current leader, all other processes receive a notification, and execute `getChildren` on `/election` to obtain the current list of children of `/election`.

If the number of clients is large, it causes a spike on the number of operations that ZooKeeper servers have to process.

To avoid the herd effect, it is sufficient to watch for the next znode down on the sequence of znodes. If a client receives a notification that the znode it is watching is gone, then it becomes the new leader in the case that there is no smaller znode. This avoids the herd effect by not having all clients watching the same znode.

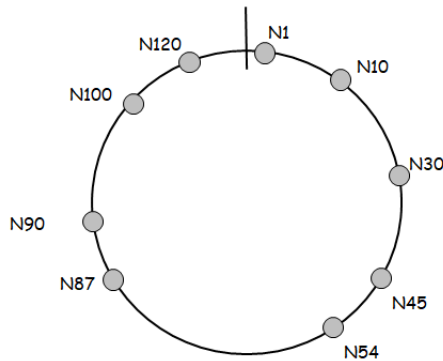
Note: Students may include a diagram to explain this.

(d) [6 marks] Trace out the steps of the Chord protocol for supporting the look-up service in a P2P network implemented using a hash space of size 2^7 .

(i) [1 marks] Draw the hash ring to accommodate a hash space of 2^7 with N1, N10, N30, N45, N54, N87, N90, N100 and N120 as the active nodes in the network.

(ii) [3 marks] Calculate the finger table for N87.

(iii) [2 marks] Illustrate the process of look-up for K26 initiated at N87.



Finger table for N87			
Hash space: $2^7 = 128$ (0 - 127)			
i	2^i	Key ID	Successor ID
0	1	$87 + 1 = 88$	90
1	2	$87 + 2 = 89$	90
2	4	$87 + 4 = 91$	100
3	8	$87 + 8 = 95$	100
4	16	$87 + 16 = 103$	120
5	32	$87 + 32 = 119$	120
6	64	$87 + 64 = 151(\text{mod } 128) = 23$	30

Look-up for K26 initiated at N87:

Node N87 will look at its finger table to identify the node which is nearest to K26: N120. It then passes the lookup operation to N120. Which then looks-up its finger table to identify the successor of K26 which is N30 and passes this information back to N87.

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.