Midterm Exam
Term 251

Course title: Distributed Systems **Course Code**: CS 435

Exam date: 27 October 2025 **Exam Time: 60** minutes

Student Name: _____

Student ID:									

Question	CLO	Max point	Score
1	1	5	
2	2	10	
3	3	5	
	Total	20	

Instructions:

- This exam contains three questions with multiple parts.
- Time allowed: 60 minutes
- Closed Book, Closed Notes.
- Use of Calculators and / or computing devices / smartphones etc is strictly prohibited.
- Answer the problems on the exam sheets only. No additional attachments would be accepted.
- If you need extra space use the back of a page.
- When the "time is over" is called, it is students' responsibility to submit his exam to the invigilator. Submitting completed exam 3 minutes after the "time is over" will incur a penalty of **5 points**.
- Do **NOT** use the erasable pens

Write your answers for MCQ type questions in this area only. DO NOT write answers anywhere else (it will not be graded).

1	Ⓐ	Ⓑ	Ⓒ	Ⓓ
2	Ⓐ	Ⓑ	Ⓒ	Ⓓ
3	Ⓐ	Ⓑ	Ⓒ	Ⓓ
4	Ⓐ	Ⓑ	Ⓒ	Ⓓ
5	Ⓐ	Ⓑ	Ⓒ	Ⓓ
6	Ⓐ	Ⓑ	Ⓒ	Ⓓ
7	Ⓐ	Ⓑ	Ⓒ	Ⓓ
8	Ⓐ	Ⓑ	Ⓒ	Ⓓ
9	Ⓐ	Ⓑ	Ⓒ	Ⓓ
10	Ⓐ	Ⓑ	Ⓒ	Ⓓ

11	Ⓐ	Ⓑ	Ⓒ	Ⓓ
12	Ⓐ	Ⓑ	Ⓒ	Ⓓ
13	Ⓐ	Ⓑ	Ⓒ	Ⓓ
14	Ⓐ	Ⓑ	Ⓒ	Ⓓ
15	Ⓐ	Ⓑ	Ⓒ	Ⓓ
16	Ⓐ	Ⓑ	Ⓒ	Ⓓ
17	Ⓐ	Ⓑ	Ⓒ	Ⓓ
18	Ⓐ	Ⓑ	Ⓒ	Ⓓ
19	Ⓐ	Ⓑ	Ⓒ	Ⓓ
20	Ⓐ	Ⓑ	Ⓒ	Ⓓ
21	Ⓐ	Ⓑ	Ⓒ	Ⓓ
22	Ⓐ	Ⓑ	Ⓒ	Ⓓ

23	Ⓐ	Ⓑ	Ⓒ	Ⓓ
24	Ⓐ	Ⓑ	Ⓒ	Ⓓ

Question 1 [5 points - CLO 1]
Part A [5]- Select the most appropriate choice from the following MCQs

1	In the context of operating systems for distributed computing, which statement is **most accurate**? A. A multi-computer OS provides independent kernels without shared memory. B. A network OS offers global process scheduling and unified file systems. C. A middleware-based OS abstracts distributed resources through standardized APIs. D. A uni-computer OS manages communication over physical networks.
2	In the "Scaling" discussion, *vertical scaling* primarily refers to which of the following? A. Adding more machines to a distributed system B. Increasing resources (CPU, memory) within an existing machine C. Replicating databases across multiple locations D. Using P2P networks to share processing load
3	Which of the following transparency types *hides resource duplication* from the user, ensuring they cannot tell whether multiple copies exist? A. Replication transparency B. Location transparency C. Migration transparency D. Access transparency
4	A startup company wants to build and deploy its own web application but doesn't want to manage the underlying servers or networking. Instead, the cloud provider offers them a complete environment including development tools, runtime, and APIs so the team can focus only on writing and running their code. Which type of cloud service are they using? A. SaaS B. PaaS C. IaaS D. DaaS
5	You are asked to build a Hadoop cluster using 40 old PCs from the IT department's storage, each with different CPU and memory configurations. Even though the machines are low-cost and heterogeneous, the system runs as one coordinated unit managed centrally by Hadoop. In this setup, what characteristic makes the cluster behave more like a high-performance computing (HPC) system rather than a grid computing system? A. It is centrally managed and runs parallel tasks on shared local resources B. It uses heterogeneous hardware and distributed ownership C. It operates primarily over the public Internet D. It requires each node to be independently scheduled and managed
6	You are designing a system to process massive datasets using multiple multi-core CPUs. Some groups of processors share a local memory within a node, while nodes communicate with each other over a network. Which statement best describes this system architecture? A. A shared-memory parallel computer. B. A distributed-memory parallel computer. C. A hybrid parallel computer. D. A concurrent system, since multiple tasks may run at the same time
7	A distributed system runs across multiple servers. Most of the time, messages between servers are delivered within a predictable time bound. However, during occasional network congestion or brief failures, message delivery may take longer than expected, and the system temporarily behaves unpredictably. Which timing model best describes this system? A. Synchronous system B. Partial-synchronous system C. Asynchronous system D. Concurrent system

8	In a *data-centric architecture*, how do producers and consumers interact with shared information? A. Producers push updates directly to consumers B. Consumers pull data from a shared repository maintained by producers C. Producers and consumers exchange messages over a peer-to-peer channel D. Both use middleware queues without centralized data
9	In a P2P system, if a peer leaves the network unexpectedly, what is the *most direct impact*? A. The network fails entirely since discovery depends on the central server. B. The departure affects only its own resources; other peers continue operating independently. C. All active connections must reset to synchronize states. D. The network must immediately elect a new central peer for coordination.
10	A company has only one webserver for their e-commerce business; they would invest to include two more web-servers. What major *distributed-system goal* does this design primarily address? A. Consistency B. Scalability C. Transparency D. Security

Question 2 [9 points - CLO 2]

Part A [6]- Select the most appropriate choice from the following MCQs

11	According to Amdahl's Law, if 25% of a program is inherently sequential, what is the maximum theoretical speedup achievable regardless of the number of processors? A. 2× B. 3× C. 4× D. 5×
12	In a Fork–Join divide-and-conquer algorithm, which operation corresponds to creating a new branch in the task DAG? A. The start() method invocation. B. The join() synchronization. C. The fork() call. D. The recursive return.
13	When implementing a divide-and-conquer algorithm for a large dataset using parallel computing, why is a fork-join approach preferred over creating fixed number of threads? A. It limits recursion depth and reduces memory usage. B. It allows only one processor to perform combining operations. C. It enforces strict sequential execution order. D. It balances workload dynamically and minimizes idle time across processors.
14	In the context of MapReduce-based distributed parallel processing, what primarily distinguishes Map and Reduce stages? A. Map performs local computation; Reduce performs synchronization. B. Map emits key-value pairs; Reduce aggregates all values for the same key. C. Map works on small data chunks; Reduce divides them into subchunks. D. Map combines results; Reduce executes parallel transformation.

15	Why does naively creating too many threads (one per small task) degrade performance, even in a fork–join model? A. The cost of context switching and scheduling overhead outweighs the computation cost. B. Threads always execute serially, defeating parallelism. C. More threads guarantee contention-free memory access. D. The JVM disables garbage collection when thread count is high.
16	A divide-and-conquer array-sum algorithm achieves `T1 = O(n)` and `T∞ = O(log n)`. Assuming a perfect scheduler and P processors, what is the expected asymptotic runtime `TP` according to the work–span model? A. O(n log n / P) B. O(n / log n + P) C. O(P + n log n) D. O(n / P + log n)
17	Referring to sockets in Java; a key architectural difference between a TCP `ServerSocket` and the connecting `Socket` in Java is that: A) ServerSocket is for UDP, while Socket is for TCP. B) ServerSocket is a passive socket for accepting connections, Socket is an endpoint for data exchange. C) ServerSocket handles data encryption, while Socket manages data compression. D) ServerSocket is used by the client to initiate a connection, while Socket is used by the server to listen.
18	A distributed application requires minimal latency, can tolerate packet loss, and must immediately send received data to the application. The most significant trade-off in choosing UDP over TCP for this application is the loss of: A) In-order, reliable byte streams and congestion control. B) The ability to organize bits into frames for hop-to-hop delivery. C) The physical transmission of bits over the network medium. D) Dialog control and synchronization between endpoints.
19	In an RPC system, what does marshalling do? A) Allocate memory on the stack for the procedure's parameters. B) Resolve the symbolic name of the procedure to a memory address. C) Transform different data representations into a byte sequence for transmission. D) Manage the CPU's program counter and register states during the call.
20	A difference between Java RMI and Sun RPC is that Java RMI can pass remote object stubs as parameters, which enables: A) The use of an Interface Definition Language (IDL) for service definition. B) The distribution of marshalling logic to a separate compiler. C) Implicit typing of parameters using a system like XDR. D) Method invocations can be made on objects that reside on different JVMs.
21	In the context of RMI failure semantics, an "orphan" is most precisely defined as a: A) Lost request message that is retransmitted by the client stub. B) Server-side computation that is executing on behalf of a client that has already crashed. C) Duplicate reply message filtered out by the client's communication module. D) Idempotent operation that can be safely executed multiple times.
22	Web Services are better alternatives to RPCs because: A) The performance characteristics and trust boundaries of Local Area Networks (LANs). B) The object-oriented programming model found in languages like Java and C++. C) The use of implicit typing systems for data representation. D) The lack of support for interface definition languages.

Part B [4]- Write **ONE line short Answer(s)** for each of the following:

- Explain why the speedup of a program is limited by its sequential part?

- Ahmed claims that his program is X% parallelizable. He plans to execute his program on the state-of-the-art 200,000 Saudi Riyals HPC server recently bought at PSU, with 40 cores. After running the program, he is able to get a speedup of 4. Estimate X.

Amdahl's Law:

$$\text{Speedup} = \frac{1}{(1-P) + \frac{P}{N}}$$

Question 3 [5 points - CLO 3]

Part A [1]- Select the most appropriate choice from the following MCQs

23	What does the Byzantine Generals Problem illustrate in distributed systems? A. Scalability B. Availability C. Consensus D. Security	
24	Which scenario most clearly demonstrates strong consistency? A. In an online banking app, a user transfers money, and immediately both devices show the updated balance. B. In a social media app, a user posts a photo, but it takes a few seconds before it appears to friends. C. In a content delivery network, cached data takes time to refresh globally. D. In an IoT system, sensor readings are stored asynchronously in multiple nodes.	

Part B [4]- A company named ShopX runs an online e-commerce platform built on a distributed system. The system has the following components spread across multiple servers:

- Frontend servers handle user requests (HTTP traffic).
- Application servers manage business logic (cart, checkout, payments).
- Database cluster (3 nodes) stores user data and transactions using a replicated setup (primary + two replicas).

One day, during a flash sale, the primary database node suddenly crashes due to hardware failure. Thousands of transactions are in progress. The system must remain available and consistent to avoid losing sales. For each of the following, explain (short answer one-line only) how would you address:

- How is failure detected?

- Explain how strong consistency will be maintained across the replicas if a user's request updates the inventory in the database.

-END OF EXAM-