



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**  
(An Autonomous Institute from A. Y. 2022-23)

**End-Sem Examination- Winter 2025  
Model Answer**

<b>Academic Year: 2025-2026</b>	<b>Semester: I</b>
<b>Name of Programme: MCA</b>	<b>Pattern: 2024</b>
<b>Name of Course: Database System and SQL</b>	<b>Course Code: 2409512</b>
<b>Max. Marks: 60</b>	<b>Duration: 2:30Hr.</b>

Q. No.	Details	Max. Marks
1	<p>Construct ER Diagram for Library Management System: A college library maintains records of books, members, issue/return transactions, and fines. Each book has a title, ISBN, author, and category. Members can borrow multiple books, and fines apply for late returns</p> <p><b>Answer:</b></p> <p>1) Entities &amp; Attributes</p> <ul style="list-style-type: none"> <li>• Book (ISBN PK, Title, Author, Category) <ul style="list-style-type: none"> <li>○ <i>Note:</i> Author may be multivalued if books can have multiple authors.</li> </ul> </li> <li>• Member (MemberID PK, Name, Address, Phone, Email, MembershipDate)</li> <li>• IssueTransaction (IssueID PK, IssueDate, DueDate, ReturnDate, FineAmount) <ul style="list-style-type: none"> <li>○ FineAmount can be stored or derived (derived = f(DueDate, ReturnDate)).</li> </ul> </li> </ul> <p>2) Relationships and Cardinalities</p> <ul style="list-style-type: none"> <li>• Issues (between Member and Book) — implemented as associative entity IssueTransaction. <ul style="list-style-type: none"> <li>○ A Member can issue 0..* IssueTransaction records (borrow many books over time).</li> <li>○ A Book can be involved in 0..* IssueTransaction records (can be borrowed multiple times over time).</li> <li>○ So cardinality between Member and IssueTransaction = 1 Member : * IssueTransaction; between Book and IssueTransaction = 1 Book : * IssueTransaction.</li> </ul> </li> <li>• Participation: IssueTransaction totally depends on both Member and Book (each transaction must reference a member and a book).</li> </ul> <p><b>3) ER Diagram (ASCII sketch)</b></p> <p>+-----+            +-----+            +-----+</p>	[6]



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<pre> Member   1 *   IssueTransaction   * 1   Book   ----- ----- ----- ----- -----  MemberID (PK)   IssueID (PK)   ISBN (PK)   Name   IssueDate   Title   Address   DueDate   Author   Phone   ReturnDate   Category   Email   FineAmount (derived)   +-----+ +-----+ +-----+ </pre> <p>Relationship: Member ---&lt; Issues (IssueTransaction) &gt;--- Book</p> <p>Notes:</p> <ul style="list-style-type: none"> <li>- IssueTransaction is an associative (weak) entity capturing the many-to-many borrow history between Member and Book.</li> <li>- FineAmount can be computed at return time as <math>\max(0, (\text{ReturnDate} - \text{DueDate}) * \text{rate})</math>.</li> </ul>	
<b>2</b>	<p>Consider the following database schema:          STUDENT(StudentID, StudentName, Program, City)          ENROLLMENT(EnrollID, StudentID, CourseID, Semester)          COURSE(CourseID, CourseName, Credits)</p> <p>Write SQL queries to perform the following:</p> <ol style="list-style-type: none"> <li>a) Display the names of students enrolled in the “Database Systems” course.</li> <li>b) Retrieve the list of distinct cities from which students belong.</li> <li>c) Display students who are enrolled in more than 3 courses</li> </ol> <p>Ans:</p> <ol style="list-style-type: none"> <li>a) Display the names of students enrolled in the “Database Systems” course.              SELECT S.StudentName              FROM STUDENT S              JOIN ENROLLMENT E ON S.StudentID = E.StudentID              JOIN COURSE C ON E.CourseID = C.CourseID              WHERE C.CourseName = 'Database Systems';</li> <li>(b) Retrieve the list of distinct cities from which students belong.              SELECT DISTINCT City              FROM STUDENT;</li> <li>(c) Display students who are enrolled in more than 3 courses.              SELECT S.StudentID, S.StudentName, COUNT(E.CourseID) AS              TotalCourses              FROM STUDENT S              JOIN ENROLLMENT E ON S.StudentID = E.StudentID              GROUP BY S.StudentID, S.StudentName              HAVING COUNT(E.CourseID) &gt; 3;</li> </ol>	[6]
<b>Q.3</b>	<p><b>a)</b> Given a relation that violates 1NF, <b>explain</b> how you would reorganize the structure to enforce atomicity and 1NF. Demonstrate your explanation using your own simple example. (8 Marks)</p>	[16]



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

Ans:

A relation violates First Normal Form (1NF) when it contains multivalued attributes, unordered sets, repeating groups, or mixed data types. To bring such a relation into 1NF, every attribute must contain only atomic (indivisible) values, and each field must store a single value of a single type.

Steps to enforce atomicity and convert to 1NF

1. Identify non-atomic attributes  
Check for fields that contain lists, sets, or multiple values (e.g., “Java, Python”, “Pen, Pencil”).
2. Split multivalued attributes  
Break composite or repeated values into separate rows or separate relations.
3. Ensure each attribute contains only one value  
Each cell must represent a single value, not a collection.
4. Create new relations if required  
If an attribute holds multiple values per entity, decompose the relation.
5. Maintain primary key integrity  
Ensure the resulting tables still have a valid primary key.

Example of Violating 1NF

Consider the following relation:

COURSE\_REGISTRATION

<b>StudentID</b>	<b>StudentName</b>	<b>CoursesRegistered</b>
101	Aarti	DBMS, Java, Networks
102	Rohan	Python, AI

The attribute CoursesRegistered contains multivalued and unordered sets → violates 1NF.

Conversion into 1NF

To enforce atomicity, we create separate rows for each course:

COURSE\_REGISTRATION\_1NF

<b>StudentID</b>	<b>StudentName</b>	<b>Course</b>
101	Aarti	DBMS
101	Aarti	Java
101	Aarti	Networks
102	Rohan	Python
102	Rohan	AI

**OR**

**b) Differentiate between BCNF and 3NF (8 Marks)**

Ans:

Third Normal Form (3NF) and Boyce–Codd Normal Form (BCNF) are



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

<p>higher-level normalization forms used to remove redundancy and update anomalies in relational database design. Although both aim to eliminate transitive and partial dependencies, BCNF is stricter than 3NF.</p> <p>1. Definition</p> <ul style="list-style-type: none"><li>• 3NF: A relation is in 3NF if it is in 2NF and every non-prime attribute is non-transitively dependent on the primary key. OR For every functional dependency <math>X \rightarrow Y</math>, either<ul style="list-style-type: none"><li>○ <math>X</math> is a superkey, or</li><li>○ <math>Y</math> is a prime attribute (part of a candidate key).</li></ul></li><li>• BCNF: A relation is in BCNF if for every functional dependency <math>X \rightarrow Y</math>, <math>X</math> must be a superkey. No exceptions allowed.</li></ul> <p>2. Strictness</p> <ul style="list-style-type: none"><li>• 3NF is less strict because it allows dependencies where the determinant is <i>not</i> a superkey if the dependent attribute is prime.</li><li>• BCNF is stricter because the determinant must <i>always</i> be a superkey.</li></ul> <p>3. Handling of Anomalies</p> <ul style="list-style-type: none"><li>• 3NF reduces most anomalies, but some anomalies may still remain if a dependency violates BCNF conditions.</li><li>• BCNF eliminates all redundancy-based anomalies, providing a more robust design.</li></ul> <p>4. Functional Dependency Condition</p> <ul style="list-style-type: none"><li>• 3NF: Allows <math>X \rightarrow Y</math> even if <math>X</math> is not a superkey, provided <math>Y</math> is a prime attribute.</li><li>• BCNF: Does not allow any dependency where <math>X</math> is not a superkey.</li></ul> <p>5. Decomposition</p> <ul style="list-style-type: none"><li>• 3NF decomposition is always possible without loss and often dependency-preserving.</li><li>• BCNF decomposition is always lossless, but may not preserve all functional dependencies.</li></ul>	
<p><b>c)</b> Given the relation: R2(ProjectID, EmployeeID, EmpName, ProjectName, HoursWorked) Functional Dependencies:</p> <ul style="list-style-type: none"><li>• ProjectID, EmployeeID <math>\rightarrow</math> HoursWorked</li><li>• EmployeeID <math>\rightarrow</math> EmpName</li><li>• ProjectID <math>\rightarrow</math> ProjectName</li></ul> <p>Candidate Key: (ProjectID, EmployeeID) Determine the highest normal form of relation R2. (8 marks)</p> <p><b>Ans:</b> Assume attributes are atomic <math>\rightarrow</math> R2 is in 1NF. 2NF requires no partial dependency of a non-prime attribute on a proper subset</p>	



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

<p>of a candidate key.</p> <ul style="list-style-type: none"> <li>• EmployeeID → EmpName — EmployeeID is a part of the key and determines non-prime attribute EmpName → partial dependency (violates 2NF).</li> <li>• ProjectID → ProjectName — ProjectID (part of key) determines ProjectName → partial dependency (violates 2NF).</li> </ul> <p>Because of these partial dependencies, R2 is NOT in 2NF.  <b>Since R2 is in 1NF but fails 2NF, the highest normal form satisfied by R2 is 1NF.</b></p>																							
<b>OR</b>																							
<p>d) Demonstrate Domain, Referential Integrities and Enterprise Constraints with suitable example.</p> <p><b>Ans:</b></p> <p>1. Domain Integrity</p> <p>Domain integrity ensures that each attribute in a relation has a valid and permissible value according to its data type, range, or format.</p> <p>Example</p> <p>Consider the relation:          STUDENT(StudentID, StudentName, Age, Email)</p> <table border="1" style="width: 100%; border-collapse: collapse; margin: 10px 0;"> <thead> <tr> <th style="text-align: left;">Domain Rule</th> <th style="text-align: left;">Attribute</th> <th style="text-align: left;">Valid Example</th> <th style="text-align: left;">Invalid Example</th> </tr> </thead> <tbody> <tr> <td>Age must be between 18–30</td> <td>Age</td> <td>21</td> <td>45 (violates domain)</td> </tr> <tr> <td>Email must contain “@”</td> <td>Email</td> <td><a href="mailto:raju@gmail.com">raju@gmail.com</a></td> <td>rajugmail.com</td> </tr> <tr> <td>StudentName must be text</td> <td>StudentName</td> <td>Priya</td> <td>123 (invalid type)</td> </tr> </tbody> </table> <p>Valid insertion:          INSERT INTO STUDENT VALUES (101, 'Anita', 20, 'anita@gmail.com');</p> <p>Invalid insertion violating domain:          INSERT INTO STUDENT VALUES (102, 'Rohan', 35, 'rohan123.com');</p> <ul style="list-style-type: none"> <li>• Age = 35 (out of allowed range)</li> <li>• Email missing “@”</li> </ul> <p>Thus, domain integrity is violated.</p> <p>2. Referential Integrity</p> <p>Referential integrity ensures that a foreign key in one table must refer to an existing primary key in another table.</p> <p>Example</p> <p>DEPARTMENT(DeptID, DeptName)          EMPLOYEE(EmpID, EmpName, DeptID)</p> <p>DEPARTMENT Table:</p> <table border="1" style="width: 100%; border-collapse: collapse; margin: 10px 0;"> <thead> <tr> <th style="text-align: left;">DeptID</th> <th style="text-align: left;">DeptName</th> </tr> </thead> <tbody> <tr> <td>10</td> <td>CS</td> </tr> <tr> <td>20</td> <td>IT</td> </tr> </tbody> </table> <p>EMPLOYEE Table:</p>		Domain Rule	Attribute	Valid Example	Invalid Example	Age must be between 18–30	Age	21	45 (violates domain)	Email must contain “@”	Email	<a href="mailto:raju@gmail.com">raju@gmail.com</a>	rajugmail.com	StudentName must be text	StudentName	Priya	123 (invalid type)	DeptID	DeptName	10	CS	20	IT
Domain Rule	Attribute	Valid Example	Invalid Example																				
Age must be between 18–30	Age	21	45 (violates domain)																				
Email must contain “@”	Email	<a href="mailto:raju@gmail.com">raju@gmail.com</a>	rajugmail.com																				
StudentName must be text	StudentName	Priya	123 (invalid type)																				
DeptID	DeptName																						
10	CS																						
20	IT																						



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<p><b>EmpID EmpName DeptID</b></p> <p>501 Meera 10 502 Rahul 30</p> <p>INSERT INTO EMPLOYEE VALUES (502, 'Rahul', 30); Since 30 is not present in DEPARTMENT, the foreign key is invalid → referential integrity violation.</p> <p>3. Enterprise Constraints Enterprise constraints (business rules) are organization-specific rules that the database must enforce. These are not defined by the DBMS structure but by business policies. Example Business Rule “An employee cannot work on more than one full-time project at the same time.”</p> <p>Tables: PROJECT_ASSIGN(EmpID, ProjectID, AssignmentType) (AssignmentType = 'Full-Time' or 'Part-Time')</p> <p>Valid Records:</p> <table border="1"> <thead> <tr> <th>EmpID</th> <th>ProjectID</th> <th>AssignmentType</th> </tr> </thead> <tbody> <tr> <td>501</td> <td>P01</td> <td>Full-Time</td> </tr> </tbody> </table> <p>Invalid Insertion: INSERT INTO PROJECT_ASSIGN VALUES (501, 'P02', 'Full-Time'); Employee 501 is already assigned to a Full-Time project → Assigning another Full-Time project violates the enterprise rule. To enforce this constraint:</p> <ul style="list-style-type: none"> <li>• Use triggers, or</li> <li>• Application-level validation.</li> </ul>	EmpID	ProjectID	AssignmentType	501	P01	Full-Time	
EmpID	ProjectID	AssignmentType						
501	P01	Full-Time						
<p><b>Q.4</b></p>	<p>a) Describe the ACID properties of a transaction. How does each property contribute to maintaining database consistency and reliability? Ans: <b>ACID Properties of a Transaction</b> A transaction in a database system is a sequence of operations performed as a single logical unit of work. The correctness and reliability of transactions are ensured by the ACID properties, namely Atomicity, Consistency, Isolation, and Durability.</p> <ol style="list-style-type: none"> <li>1. <b>Atomicity</b> Atomicity ensures that a transaction is treated as an indivisible unit. Either all the operations of the transaction are successfully executed, or none of them are applied to the database. If any failure occurs during transaction execution, the database is rolled back to its previous consistent state. This prevents partial updates and helps maintain database reliability.</li> <li>2. <b>Consistency</b> Consistency ensures that a transaction brings the database from one</li> </ol>	<p>[16]</p>						



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

valid (consistent) state to another valid state. All integrity constraints, rules, and triggers defined on the database must be satisfied before and after the execution of a transaction. This property guarantees that the database never violates its integrity constraints.

**3. Isolation**

Isolation ensures that the concurrent execution of multiple transactions does not interfere with each other. Each transaction is executed as if it is the only transaction in the system. Intermediate results of one transaction are not visible to others until the transaction is committed, thereby preventing anomalies such as dirty reads, lost updates, and uncommitted data access.

**4. Durability**

Durability ensures that once a transaction is successfully committed, its effects are permanently recorded in the database, even in the presence of system failures such as power outages or crashes. This is typically achieved through logging and recovery mechanisms, ensuring long-term reliability of the database.

**OR**

b) Explain the concept of lock-based concurrency control. Why are locks required to manage concurrent transactions in a database system?

Ans:

**Lock-Based Concurrency Control**

Lock-based concurrency control is a technique used in database management systems to control the concurrent execution of transactions by regulating access to data items through **locks**. Before a transaction can read or write a data item, it must first obtain an appropriate lock on that item. This ensures that only authorized transactions can access the data at a given time.

Locks are mainly of two types:

- **Shared (S) lock:** Allows a transaction to read a data item. Multiple transactions can hold shared locks on the same data item simultaneously.
- **Exclusive (X) lock:** Allows a transaction to read and write a data item. Only one transaction can hold an exclusive lock on a data item at a time.

**Why Locks Are Required in Concurrent Transactions**

Locks are required to manage concurrent transactions for the following reasons:

**1. To prevent data inconsistency**

Without locks, multiple transactions may simultaneously read or write the same data item, leading to problems such as lost updates, dirty reads, and uncommitted data access.

**2. To maintain isolation**

Locks ensure that transactions execute as if they are isolated from each



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<p>other. Intermediate results of one transaction are not visible to other transactions until the transaction is completed.</p> <ol style="list-style-type: none"><li><b>To ensure serializability</b> Lock-based protocols help ensure that the effect of concurrent transactions is equivalent to some serial execution, thereby preserving database correctness.</li><li><b>To handle conflicting operations</b> Locks control access to data items when read–write or write–write conflicts occur, ensuring correct execution order.</li><li><b>To support ACID properties</b> By enforcing isolation and consistency, lock-based concurrency control plays a key role in maintaining the ACID properties of transactions.</li></ol>	
	<p>c) Explain what is meant by a transaction schedule. Differentiate between serial and non-serial schedules with suitable explanation.</p> <p>Ans:</p> <p><b>Transaction Schedule</b> A <b>transaction schedule</b> is the chronological order in which the operations (read, write, commit, abort) of one or more transactions are executed by the database management system. It shows how multiple transactions are interleaved when they are executed concurrently, while preserving the order of operations within each individual transaction.</p> <p><b>Serial Schedule</b> A <b>serial schedule</b> is one in which transactions are executed one after another, without any interleaving of operations. One transaction completes all its operations before the next transaction starts.</p> <p><b>Explanation:</b> If there are two transactions T1 and T2, in a serial schedule either:</p> <ul style="list-style-type: none"><li>• T1 executes completely first, followed by T2, or</li><li>• T2 executes completely first, followed by T1.</li></ul> <p><b>Example:</b> T1: Read(A) → Write(A) → Commit T2: Read(B) → Write(B) → Commit Here, T2 starts only after T1 finishes.</p> <p><b>Characteristics:</b></p> <ul style="list-style-type: none"><li>• Simple and easy to understand</li><li>• Always maintains database consistency</li><li>• Low concurrency and poor performance</li></ul> <p><b>Non-Serial Schedule</b> A <b>non-serial schedule</b> is one in which the operations of multiple transactions are interleaved. Transactions execute concurrently, but the internal order of operations of each transaction is preserved.</p> <p><b>Explanation:</b></p>	



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<p>Operations of T1 and T2 are mixed to improve system performance and resource utilization.</p> <p><b>Example:</b>  T1: Read(A)  T2: Read(B)  T1: Write(A)  T2: Write(B)  T1: Commit  T2: Commit</p> <p><b>Characteristics:</b></p> <ul style="list-style-type: none"> <li>• Allows concurrent execution of transactions</li> <li>• Improves system throughput and resource utilization</li> <li>• May cause inconsistency if not properly controlled</li> <li>• Requires concurrency control techniques</li> </ul> <p><b>Difference Between Serial and Non-Serial Schedules</b></p> <table border="1"> <thead> <tr> <th>Aspect</th> <th>Serial Schedule</th> <th>Non-Serial Schedule</th> </tr> </thead> <tbody> <tr> <td>Execution</td> <td>One transaction at a time</td> <td>Transactions execute concurrently</td> </tr> <tr> <td>Interleaving</td> <td>No interleaving</td> <td>Interleaving of operations</td> </tr> <tr> <td>Consistency</td> <td>Always consistent</td> <td>Consistent only if serializable</td> </tr> <tr> <td>Performance</td> <td>Low</td> <td>High</td> </tr> <tr> <td>Complexity</td> <td>Simple</td> <td>Requires concurrency control</td> </tr> </tbody> </table> <p><b>OR</b></p> <p>d) Explain what a deadlock is in a database system. How does deadlock handling ensure smooth execution of transactions?</p> <p>Ans:  A deadlock in a database system occurs when two or more transactions are waiting indefinitely for each other to release locks on data items. In this situation, none of the transactions can proceed because each transaction holds a lock that the other transaction needs, resulting in a circular wait condition.</p> <p><b>Example:</b></p> <ul style="list-style-type: none"> <li>• Transaction T1 holds a lock on data item A and waits for a lock on data item B.</li> <li>• Transaction T2 holds a lock on data item B and waits for a lock on data item A.</li> </ul> <p>This circular waiting leads to a deadlock.</p> <p><b>Deadlock Handling and Smooth Transaction Execution</b>  Deadlock handling ensures smooth execution of transactions by identifying and resolving deadlock situations so that the database system does not remain blocked. This is achieved through the following approaches:</p> <ol style="list-style-type: none"> <li>1. <b>Deadlock Prevention</b>  Deadlock prevention techniques avoid deadlocks by ensuring that at</li> </ol>	Aspect	Serial Schedule	Non-Serial Schedule	Execution	One transaction at a time	Transactions execute concurrently	Interleaving	No interleaving	Interleaving of operations	Consistency	Always consistent	Consistent only if serializable	Performance	Low	High	Complexity	Simple	Requires concurrency control	
Aspect	Serial Schedule	Non-Serial Schedule																		
Execution	One transaction at a time	Transactions execute concurrently																		
Interleaving	No interleaving	Interleaving of operations																		
Consistency	Always consistent	Consistent only if serializable																		
Performance	Low	High																		
Complexity	Simple	Requires concurrency control																		



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<p>least one of the necessary conditions for deadlock (such as circular wait) never occurs. This is done using methods like timestamp-based ordering (wait-die and wound-wait).</p> <p>2. <b>Deadlock Detection</b> In deadlock detection, the system periodically checks for deadlocks using structures like a wait-for graph. If a cycle is detected, it indicates a deadlock.</p> <p>3. <b>Deadlock Recovery</b> Once a deadlock is detected, the system resolves it by aborting one or more transactions (victims) involved in the deadlock and releasing their locks so that other transactions can continue.</p>	
<p style="text-align: center;"><b>Q.5</b></p>	<p>a) Analyze the differences between centralized database architecture and client-server database architecture in terms of data access, performance, scalability, and fault tolerance.</p> <p>Ans: Centralized and client-server database architectures differ significantly in how data is accessed, how the system performs under load, how it scales, and how it handles failures. The key differences are analyzed below:</p> <p>1. Data Access</p> <p>In a centralized database architecture, all data and the database management system reside on a single central machine. Users access data through terminals or thin clients, and all processing is performed at the central server. This creates a single point of data access.</p> <p>In a client-server database architecture, the database server stores and manages the data, while multiple client machines handle user interfaces and some application logic. Clients send queries to the server, which processes them and returns results. This distributed access allows better utilization of system resources.</p> <p>Analysis: Client-server architecture provides more flexible and efficient data access by distributing workload between clients and server, whereas centralized systems rely entirely on one machine.</p> <p>2. Performance</p> <p>In a centralized architecture, performance may degrade as the number of users increases because all requests are processed by a single system. Resource contention for CPU, memory, and I/O can become a bottleneck.</p> <p>In a client-server architecture, performance is generally better since client machines handle presentation and partial processing, reducing the server's</p>	<p style="text-align: center;">[16]</p>



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

<p>workload. Concurrent processing and better resource utilization improve overall throughput.</p> <p>Analysis: Client–server systems offer higher performance for multi-user environments, while centralized systems are suitable only for small-scale applications.</p> <p>3. Scalability</p> <p>Centralized systems have limited scalability. Scaling often requires upgrading the central server (vertical scaling), which can be costly and has physical limits.</p> <p>Client–server systems are more scalable. Additional clients or servers can be added to handle increased load (horizontal scaling), making it suitable for growing organizations.</p> <p>Analysis: Client–server architecture supports better scalability due to its distributed nature, whereas centralized architecture is rigid and difficult to scale.</p> <p>4. Fault Tolerance</p> <p>In a centralized architecture, failure of the central server results in complete system downtime, as it is a single point of failure.</p> <p>In a client–server architecture, failure of a client does not affect other clients. Although server failure impacts the system, backup servers and replication techniques can be used to improve fault tolerance.</p> <p>Analysis: Client–server architecture provides better fault tolerance and system availability compared to centralized architecture.</p> <p>OR</p> <p>b) Analyze the need for parallel databases. How does parallel database architecture improve query performance and throughput compared to traditional database systems?</p> <p>Ans: With the rapid growth of data volume and the increasing number of users, traditional (single-processor) database systems often fail to meet performance and scalability requirements. Parallel databases are designed to overcome these limitations by using multiple processors, disks, and memory units to execute database operations concurrently.</p> <p>Need for Parallel Databases</p> <p>1. Handling Large Data Volumes</p>	
--	--



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<p>Modern applications generate massive amounts of data. Traditional databases process queries sequentially, leading to long response times. Parallel databases divide data across multiple disks and nodes, enabling faster data access.</p> <ol style="list-style-type: none"><li>2. Improved Query Response Time Complex queries involving joins, aggregations, and scans are time-consuming in traditional systems. Parallel databases execute different parts of a query simultaneously, significantly reducing execution time.</li><li>3. Higher Transaction Throughput In multi-user environments, traditional systems become bottlenecks due to limited processing resources. Parallel databases allow multiple transactions to be processed concurrently, increasing overall system throughput.</li><li>4. Scalability Parallel databases support scalable architectures where additional processors or nodes can be added to handle increased workload, unlike traditional systems that rely mainly on vertical scaling.</li></ol> <p>How Parallel Database Architecture Improves Performance and Throughput</p> <ol style="list-style-type: none"><li>1. Data Partitioning Data is divided (horizontally or vertically) across multiple disks or nodes. Parallel processing allows simultaneous access to these partitions, reducing I/O time compared to accessing a single storage unit.</li><li>2. Parallel Query Execution Query operations such as selection, projection, join, and aggregation are executed in parallel on different processors. Each processor works on a subset of data, and results are combined, improving query performance.</li><li>3. Inter-Query and Intra-Query Parallelism<ul style="list-style-type: none"><li>o <i>Inter-query parallelism</i> allows multiple queries to execute at the same time.</li><li>o <i>Intra-query parallelism</i> allows a single query to be divided into multiple tasks and executed concurrently.</li></ul>Both types increase system efficiency and throughput.</li><li>4. Reduced Resource Contention By distributing workload across multiple processors and disks, parallel databases minimize contention for CPU, memory, and I/O resources, leading to better utilization.</li></ol> <p>Comparison with Traditional Database Systems</p> <p>Traditional databases process tasks sequentially on a single processor, resulting in limited performance and scalability. In contrast, parallel database architectures leverage concurrent processing and distributed storage, leading to faster query execution, higher throughput, and better scalability.</p>	
	<p>c) Analyze the need for distributed databases over centralized database systems. How do distributed databases address data availability and performance requirements? (8 marks)</p>	



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

Ans:

Centralized database systems store all data at a single location and process all requests on one central server. While simple to manage, such systems face serious limitations when used in large-scale, geographically distributed, or high-availability applications. Distributed databases overcome these limitations by distributing data across multiple sites connected through a network.

Need for Distributed Databases over Centralized Systems

1. **Geographical Distribution of Users**  
Modern organizations operate across multiple locations. A centralized database causes high access latency for remote users, whereas distributed databases place data closer to users, reducing response time.
2. **Improved Data Availability**  
Centralized systems suffer from a single point of failure. If the central server fails, the entire system becomes unavailable. Distributed databases replicate data across multiple sites, ensuring continued availability even if one site fails.
3. **Scalability Requirements**  
Centralized systems have limited scalability and often rely on expensive hardware upgrades. Distributed databases allow horizontal scaling by adding new sites or nodes as data and workload increase.
4. **Local Autonomy and Reliability**  
Distributed databases allow each site to manage its local data independently while still participating in a global database, improving system reliability and administrative flexibility.

How Distributed Databases Address Availability and Performance

1. **Data Replication**  
Copies of data are stored at multiple sites. If one site becomes unavailable due to failure, another site can continue serving requests, thereby improving fault tolerance and data availability.
2. **Data Fragmentation and Allocation**  
Data is divided into fragments and stored at different sites based on usage patterns. This enables faster local access and reduces unnecessary data transfer over the network.
3. **Parallel Query Processing**  
Queries can be executed simultaneously at multiple sites. Each site processes its local data fragment, and the results are combined, leading to faster query execution.
4. **Load Balancing**  
Workload is distributed across multiple nodes, preventing any single node from becoming a performance bottleneck and improving overall system throughput.

Comparison with Centralized Systems

In centralized databases, all requests compete for the same resources, leading to performance degradation and downtime risks. Distributed databases, by contrast, eliminate single points of failure, reduce access latency, and provide



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

higher throughput through parallelism and data distribution.

**OR**

d) Analyze the top-down and bottom-up distributed database design processes. How does each approach affect data fragmentation, allocation, and system scalability?

Ans:

Distributed database design defines how data is fragmented, allocated, and managed across multiple sites. Two major design approaches are top-down and bottom-up, each affecting data organization and system scalability in different ways.

**1. Top-Down Design Approach**

In the top-down approach, the design starts from a global perspective. A global conceptual schema is first created based on overall system requirements, and then it is progressively refined for distribution.

Process

- Define global conceptual schema
- Identify user views and access patterns
- Perform data fragmentation (horizontal, vertical, or hybrid)
- Allocate fragments to different sites
- Define replication and local schemas

Impact on Design Factors

- Data Fragmentation:  
Fragmentation is systematic and well-planned, based on global access requirements. This leads to optimal fragmentation that minimizes remote data access.
- Data Allocation:  
Allocation decisions consider global workload, communication cost, and data locality, resulting in efficient placement of fragments.
- System Scalability:  
The approach supports high scalability because data distribution is planned in advance. However, redesign may be required if new sites or major requirements are added later.

Strengths

- Better global optimization
- High data consistency
- Suitable for new distributed database systems

**2. Bottom-Up Design Approach**

In the bottom-up approach, design starts from existing local databases at different sites, which are then integrated into a global distributed database.

Process

- Identify existing local schemas
- Integrate schemas into a global schema
- Resolve schema conflicts (naming, structure, data types)



**K. K. Wagh Institute of Engineering Education and Research,  
Nashik**

(An Autonomous Institute from A. Y. 2022-23)

	<ul style="list-style-type: none"> <li>• Define global views and mappings</li> </ul> <p>Impact on Design Factors</p> <ul style="list-style-type: none"> <li>• Data Fragmentation: Fragmentation already exists implicitly due to separate local databases. It may not be optimal and often reflects local needs rather than global efficiency.</li> <li>• Data Allocation: Allocation is fixed based on existing data locations, offering limited flexibility for redistribution.</li> <li>• System Scalability: Scalability is easier when adding new sites, as new local databases can be integrated without redesigning the entire system. However, performance optimization may be limited.</li> </ul> <p>Strengths</p> <ul style="list-style-type: none"> <li>• Suitable for integrating legacy systems</li> <li>• Less initial redesign effort</li> <li>• Supports organizational autonomy</li> </ul> <p>3. Comparative Analysis</p> <table border="0" style="width: 100%;"> <thead> <tr> <th style="text-align: left;">Aspect</th> <th style="text-align: left;">Top-Down Approach</th> <th style="text-align: left;">Bottom-Up Approach</th> </tr> </thead> <tbody> <tr> <td>Design Start</td> <td>Global schema</td> <td>Local schemas</td> </tr> <tr> <td>Fragmentatio</td> <td>Planned and optimized</td> <td>Pre-existing, less optimized</td> </tr> <tr> <td>Allocation</td> <td>Flexible and workload-based</td> <td>Fixed, location-dependent</td> </tr> <tr> <td>Scalability</td> <td>High, but redesign may be needed</td> <td>Easy site addition, limited optimization</td> </tr> <tr> <td>Suitability</td> <td>New systems</td> <td>Existing/legacy systems</td> </tr> </tbody> </table>	Aspect	Top-Down Approach	Bottom-Up Approach	Design Start	Global schema	Local schemas	Fragmentatio	Planned and optimized	Pre-existing, less optimized	Allocation	Flexible and workload-based	Fixed, location-dependent	Scalability	High, but redesign may be needed	Easy site addition, limited optimization	Suitability	New systems	Existing/legacy systems	
Aspect	Top-Down Approach	Bottom-Up Approach																		
Design Start	Global schema	Local schemas																		
Fragmentatio	Planned and optimized	Pre-existing, less optimized																		
Allocation	Flexible and workload-based	Fixed, location-dependent																		
Scalability	High, but redesign may be needed	Easy site addition, limited optimization																		
Suitability	New systems	Existing/legacy systems																		