



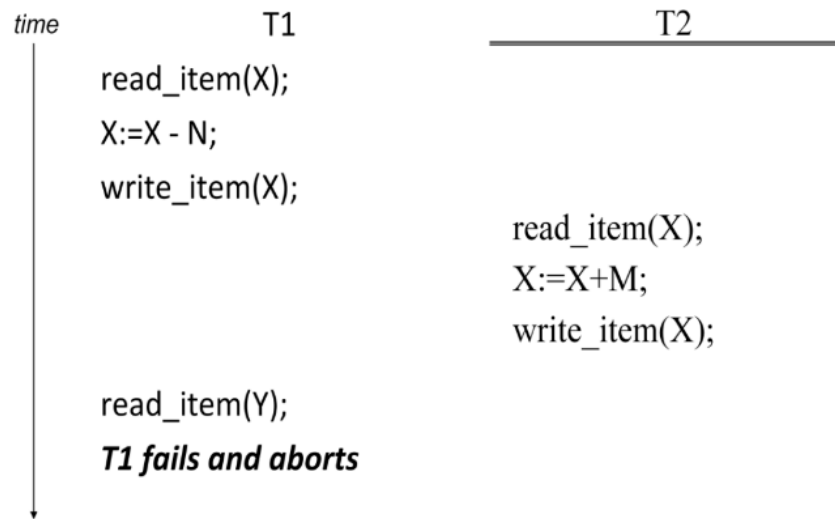
USN 

--	--	--	--	--	--	--	--	--	--




Internal Assessment Test III –Aug 2024

Sub:	Database Management System				Sub Code:	BCS/BAD 403	Branch:	AINDS / CS (DS)																																					
Date:	11/07/2024	Duration:	90 minutes	Max Marks:	50	Sem	IV	OBE																																					
<b>Answer any FIVE Questions</b>								<b>MARKS</b>	<b>CO</b>	<b>RBT</b>																																			
1	A	<p><b>1a. Describe the database inconsistency problems: Lost update, dirty read and blind write?</b></p> <p>DBMS has a <i>Concurrency Control</i> sub-system to assure database remains in consistent state despite concurrent execution of transactions</p> <p>Three problems are occurs during concurrency</p> <ol style="list-style-type: none"> <li>1. The lost update problem</li> <li>2. The temporary update (dirty read) problem</li> <li>3. Incorrect summary problem</li> </ol> <p><b>Lost Update Problem</b></p> <table style="width: 100%; border: none;"> <tr> <td style="text-align: center; width: 15%;"><i>time</i></td> <td style="width: 40%;"></td> <td style="text-align: center; width: 15%;">T1</td> <td style="width: 30%;"></td> <td style="text-align: center;">T2</td> </tr> <tr> <td></td> <td>read_item(X);</td> <td></td> <td></td> <td>read_item(X);</td> </tr> <tr> <td></td> <td>X:=X - N;</td> <td></td> <td></td> <td>X:=X+M;</td> </tr> <tr> <td></td> <td>write_item(X);</td> <td></td> <td></td> <td>write_item(X);</td> </tr> <tr> <td></td> <td>read_item(Y);</td> <td></td> <td></td> <td></td> </tr> <tr> <td></td> <td>Y:=Y + N;</td> <td></td> <td></td> <td></td> </tr> <tr> <td></td> <td>write_item(Y);</td> <td></td> <td></td> <td></td> </tr> </table> <p><b>Temporary Update (Dirty Read)</b></p>						<i>time</i>		T1		T2		read_item(X);			read_item(X);		X:=X - N;			X:=X+M;		write_item(X);			write_item(X);		read_item(Y);					Y:=Y + N;					write_item(Y);				5	CO1	L1
<i>time</i>		T1		T2																																									
	read_item(X);			read_item(X);																																									
	X:=X - N;			X:=X+M;																																									
	write_item(X);			write_item(X);																																									
	read_item(Y);																																												
	Y:=Y + N;																																												
	write_item(Y);																																												



**Incorrect Summary Problem**



B

5

CO1

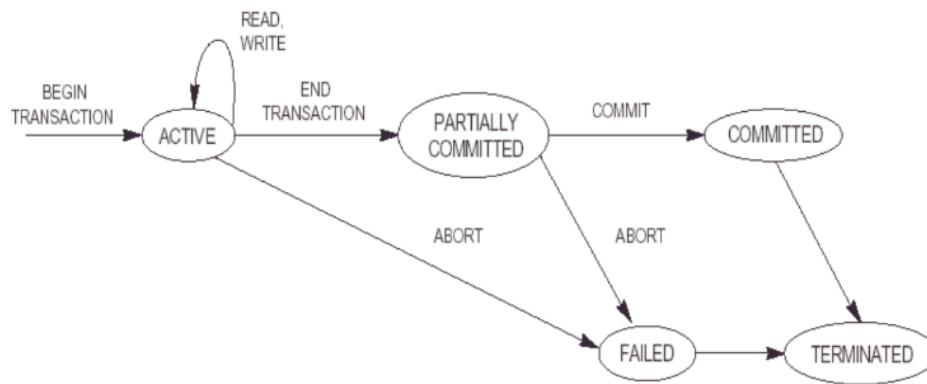
L2

**With a neat diagram, explain the various states of a transaction execution?**

**Transaction states**

- BEGIN\_TRANSACTION: marks start of transaction
- READ or WRITE: two possible operations on the data
- END\_TRANSACTION: marks the end of the read or write operations; start checking whether everything went according to plan
- COMMIT\_TRANSACTION: signals successful end of transaction; changes can be “committed” to DB
- Partially committed

- ROLLBACK (or ABORT): signals unsuccessful end of transaction, changes applied to DB must be undone



**Check whether the below schedule is conflict or not.**

**{b2 , r2(X) , b1 , r1(X) , w1(X), r1(Y), w1(Y), w2(X), e1, c1, e2, c2).**

T1	T2
	B2 R2(X)
B1 R1(X)	
W1(X) R1(Y) W1(Y)	
	W2(X)
E1 C1	
	E2 C2

The above schedule is not serializable

Before W2(X) transaction T1 is reading the R1(X) value

The serializable schedule as follows  
{b2 , r2(X) , b1 , **w2(X)**, r1(X) , **w1(X)**, r1(Y), w1(Y),, e1, c1, e2, c2}.

**What is 2PL? Explain with an example.**

2 a,b

5

CO1

L2

## Two-Phase Locking (2PL) Protocol

Transaction is said to follow the *two-phase-locking protocol* if all locking operations precede the *first* unlock operation

Expanding (growing) phase

Shrinking phase

During the shrinking phase no new locks can be acquired

Downgrading ok

Upgrading is not

### 2PL Example

T1'	T2'
read_lock(Y);	read_lock(X);
read_item(Y);	read_item(X);
write_lock(X);	write_lock(Y);
unlock(Y);	unlock(X);
read_item(X);	read_item(Y);
X:=X+Y;	Y:=X+Y;
write_item(X);	write_item(Y);
unlock(X);	unlock(Y);

Both T1' and T2' follow the 2PL protocol

Any schedule including T1' and T2' is guaranteed to be serializable

Limits the amount of concurrency

Two-phase locking protocol (2PL)

All lock operations precede the first unlock operation

Expanding phase and shrinking phase

Upgrading of locks must be done in expanding phase and downgrading of locks

must be done in shrinking phase

If every transaction in a schedule follows 2PL protocol then the schedule is guaranteed

to be serializable.

Variants of 2PL

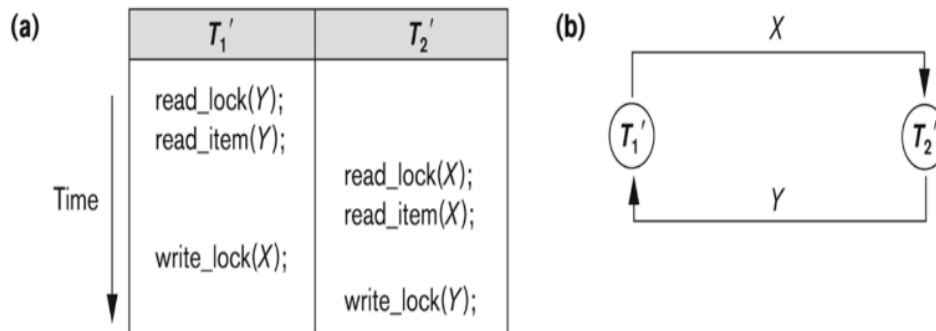
Basic, conservative, strict, and rigorous

## How do you detect a deadlock during concurrent transaction execution? Explain Deadlocks and Starvation in 2PL.

2PL can produce deadlocks

Deadlock and starvation in 2PL

Deadlock occurs when each transaction  $T$  in a set of two or more transactions is waiting for some item that is locked by some other transaction  $T'$  in the set.



Illustrating the deadlock problem. (a) A partial schedule of  $T_1'$  and  $T_2'$  that is in a state of deadlock. (b) A wait-for graph for the partial schedule in (a).

Deadlock prevention protocols

- Conservative 2PL, lock all needed items in advance

- Ordering all items in the database

Possible actions if a transaction is involved in a possible deadlock situation

- Block and wait
- Abort and restart
- Preempt and abort another transaction

Two schemes that prevent deadlock (Timestamp based)

### Wait-die

An older transaction is allowed to wait on a younger transaction whereas a younger transaction requesting an item from held by an older transaction is aborted and restarted with the same timestamp

### Wound-wait

A younger transaction is allowed to wait on an older transaction whereas an older transaction requesting an item from held by a younger transaction preempts the younger transaction by aborting it.

### Starvation

A transaction cannot proceed for an infinite period of time while other transactions in the system continue normally

- Unfair waiting scheme
- Victim selection

5

CO2

L3

3

**c) Explain the various database recovery techniques, with examples.**

- Keeps information about operations made by transactions:
  - *Before-image (undo entry)* of updated items
  - *After-image (redo entry)* of updated items
- Enables restoring a consistent state after non-catastrophic failure (forward/backward).
- Alternatives:
  - *undo/no-redo*
  - *no-undo/redo*
  - *undo/redo*
  - *no-undo/no-redo.*

**Write-Ahead Logging (WAL)**

(1) No overwrite of disk data before *undo*-type log records are forced to disk.

(2) *Both undo- and redo-type* log records (= before- and after-images) must be forced to disk before end of commit.

**Backup**

Copy of database on archival storage (off-line, often on tape).  
 Enables partial recovery from *catastrophic* failures:  
 For *committed* transactions:  
 Load backup and apply redo operations from the log (if the log survived).  
*Non-committed* transactions must be restarted (= re-executed).

**Cache**

In-memory buffer for database pages.  
 A *directory* (page table) keeps track of pages in cache.  
 Page-replacement strategy needed, e.g.  
     *FIFO* (First-In-First-Out), or  
     *LRU* (Least Recently Used)  
*Dirty bit* tells for each page, if it has changed  
*Flushing* means (force-)writing buffer pages to disk.  
*Pin/unpin bit* tells if the page can be written

**Rollback**

At failure, apply *undo*-type log records (before-images) to updated items.  
 A recoverable schedule may allow *cascading rollback*.  
 Most practical protocols avoid cascading rollbacks. Then no *read-entries* are required in the log (no dirty reads).

10

CO2

L4

**Explain binary locks and shared locks with algorithms?**



Binary locks and shared locks are two important concepts in concurrent programming and database systems. Let's explore each of them:

### 1. Binary Locks (Exclusive Locks):

Binary locks, also known as exclusive locks or mutex locks, allow only one process or thread to access a shared resource at a time. When a process acquires a binary lock, it has exclusive access to the resource, and all other processes must wait until the lock is released.

Algorithm for Binary Lock:

```
struct BinaryLock {
    bool locked = false;
}

function acquire(BinaryLock lock) {
    while (true) {
        if (!lock.locked) {
            lock.locked = true;
            return;
        }
        // Wait or yield to other processes
    }
}

function release(BinaryLock lock) {
    lock.locked = false;
}
```

### 2. Shared Locks (Read-Write Locks):

Shared locks allow multiple processes to read a shared resource simultaneously, but ensure exclusive access for writing. This is useful when read operations are more frequent than write operations.

There are two types of locks in this system:

- Read lock: Multiple processes can hold read locks simultaneously.
- Write lock: Only one process can hold a write lock, and no read locks can be held simultaneously.

Algorithm for Shared Lock:

```
struct SharedLock {
    int readers = 0;
    bool writer = false;
```

```
BinaryLock readLock;
BinaryLock writeLock;
}

function acquireReadLock(SharedLock lock) {
    acquire(lock.readLock);
    if (lock.readers == 0) {
        acquire(lock.writeLock);
    }
    lock.readers++;
    release(lock.readLock);
}

function releaseReadLock(SharedLock lock) {
    acquire(lock.readLock);
    lock.readers--;
    if (lock.readers == 0) {
        release(lock.writeLock);
    }
    release(lock.readLock);
}

function acquireWriteLock(SharedLock lock) {
    acquire(lock.writeLock);
}

function releaseWriteLock(SharedLock lock) {
    release(lock.writeLock);
}
```

**What is the CAP Theorem? Which of the three properties (consistency, availability, partition tolerance) are most important in NOSQL systems?**

The CAP Theorem, also known as Brewer's Theorem, is a fundamental concept in distributed computing systems, particularly relevant to distributed databases and NoSQL systems. It states that in a distributed data store, it is impossible to simultaneously guarantee all three of the following properties:

1. Consistency (C): All nodes see the same data at the same time. In other words, a read operation will return the most recent write operation's value, no matter which node it contacts.

	<p>2. Availability (A): Every request receives a response, without guarantee that it contains the most recent version of the data. The system remains operational and can respond to requests even if some nodes are down.</p> <p>3. Partition Tolerance (P): The system continues to operate despite network partitions or communication breakdowns between nodes.</p> <p>The theorem asserts that a distributed system can only guarantee two out of these three properties at any given time.</p> <p>In NoSQL systems, which are often designed for large-scale, distributed environments, Partition Tolerance is generally considered non-negotiable. Network partitions are a reality in distributed systems, and the ability to handle them is crucial. Therefore, NoSQL databases typically have to choose between Consistency and Availability when network partitions occur.</p> <p>Most NoSQL systems prioritize Availability and Partition Tolerance (AP) over strict Consistency. Here's why:</p> <ol style="list-style-type: none"><li>1. Scale: NoSQL databases are often used in scenarios requiring high scalability. Prioritizing availability allows the system to continue functioning even when some nodes are unreachable.</li><li>2. Performance: Strict consistency can introduce latency, as the system needs to synchronize all nodes before responding to requests. Many NoSQL use cases prioritize low-latency responses over perfect consistency.</li><li>3. Use Case Requirements: Many applications using NoSQL can tolerate eventual consistency, where the system will become consistent over time, rather than requiring immediate consistency.</li><li>4. Geographic Distribution: For globally distributed databases, network partitions are more common, and maintaining strict consistency across all regions can be challenging and performance-intensive.</li></ol>		
--	--	--	--

However, it's important to note that this is a generalization. Different NoSQL databases make different trade-offs:

- Cassandra and DynamoDB, for instance, typically favor Availability and Partition Tolerance (AP).

- MongoDB and HBase lean more towards Consistency and Partition Tolerance (CP) in their default configurations.

- Some modern distributed databases like Google's Spanner aim to provide all three properties most of the time, only sacrificing availability during rare network partitions.

The choice between consistency and availability often depends on the specific use case. For example, a financial system might prioritize consistency over availability, while a content delivery system might prioritize availability over strict consistency.

In summary, while Partition Tolerance is crucial for NoSQL systems, the choice between Consistency and Availability is often application-specific. However, many NoSQL systems are designed with a bias towards Availability and Partition Tolerance, with various mechanisms to provide different levels of eventual consistency.

--	--	--	--	--	--

**Normalize the below relation up to 3NF**

Module	Dept	Lecturer	Text
M1	D1	L1	T1
M1	D1	L1	T2
M2	D1	L1	T1
M2	D1	L1	T3
M3	D1	L2	T4
M4	D2	L3	T1
M4	D2	L3	T5
M5	D2	L4	T6

- To normalize the given relation up to Third Normal Form (3NF), we'll follow the normalization process:

**Step1:-Identify the functional dependencies.**

**Step2:-Ensure the relation is in First Normal Form (1NF).**

**Step3:-Transform it to Second Normal Form (2NF).**

**Step4:-Transform it to Third Normal Form (3NF)**

**Step 1: Identify Functional Dependencies**

- From the given table:
- Module** determines **Dept** and **Lecturer**.
- Module** and **Text** are combined to determine all attributes uniquely, as each combination of **Module** and **Text** uniquely identifies a record.
- Therefore, we have the following functional dependencies:

**Module** → **Dept, Lecturer**

**Module, Text** → **Dept, Lecturer, Text**

**Step 3: Second Normal Form (2NF)**

- A table is in 2NF if it is in 1NF and all non-key attributes are fully functionally dependent on the primary key. We need to ensure that there are no partial dependencies on a composite key.
- Current Candidate Key:** (Module, Text)
- To move to 2NF, we decompose the table to remove partial dependencies.
- Decomposition into 2NF:**

**Table 1: Modules**

Module	Dept	Lecturer
M1	D1	L1
M2	D1	L1
M3	D1	L2
M4	D2	L3
M5	D2	L4

**Table 2: Module\_Text**

Module	Text
M1	T1
M1	T2
M2	T1
M2	T3
M3	T4
M4	T1
M4	T5
M5	T6

**Step 4: Third Normal Form (3NF)**

- A table is in 3NF if it is in 2NF and all the attributes are functionally dependent only on the primary key and there are no transitive dependencies.
- Table 1: Modules** is already in 3NF since all non-key attributes (Dept, Lecturer) depend only on the primary key (Module).
- Table 2: Module\_Text** is in 3NF since there are no transitive dependencies.

**Final Normalized Tables:**

**Table 1: Modules**

Module	Dept	Lecturer
M1	D1	L1
M2	D1	L1
M3	D1	L2
M4	D2	L3
M5	D2	L4

**Table 2: Module\_Text**

Module	Text
M1	T1
M1	T2
M2	T1
M2	T3
M3	T4
M4	T1
M4	T5
M5	T6

These two tables are now in 3NF, ensuring no redundancy and eliminating partial and transitive dependencies.

--	--	--	--	--	--

5	<p><b>Demonstrate the following constraints in SQL with suitable example:</b>  <b>i) NOT NULL    ii) Primary key    iii) Foreign key    iv) Default    v) Unique</b></p> <p><b>i) NOT NULL :- Ensures that a column cannot have a NULL value</b></p> <p><b>ii) Primary key :-</b> A combination of a NOT NULL and UNIQUE. Uniquely identifies each row in a table.</p> <ul style="list-style-type: none"> <li>• The PRIMARY KEY constraint uniquely identifies each record in a table.</li> <li>• Primary keys must contain UNIQUE values, and cannot contain NULL values.</li> <li>• A table can have only ONE primary key; and in the table, this primary key can consist of single or multiple columns.</li> </ul> <p><b>SQL PRIMARY KEY on CREATE TABLE</b></p> <ul style="list-style-type: none"> <li>• The following SQL creates a PRIMARY KEY on the "ID" column when the "Persons" table is created: <pre>CREATE TABLE Persons (     ID int NOT NULL,     LastName varchar(255) NOT NULL,     FirstName varchar(255),     Age int,     PRIMARY KEY (ID));</pre> </li> </ul> <p><b>SQL PRIMARY KEY on ALTER TABLE</b></p> <p>To create a PRIMARY KEY constraint on the "ID" column when the table is already created, use the following SQL:</p> <pre>ALTER TABLE Persons ADD PRIMARY KEY (ID);</pre> <p><b>DROP a PRIMARY KEY Constraint</b></p> <p>To drop a PRIMARY KEY constraint, use the following SQL:</p> <pre>ALTER TABLE Persons DROP PRIMARY KEY;</pre> <p><b>iii) Foreign key:-</b> Prevents actions that would destroy links between tables.</p> <ul style="list-style-type: none"> <li>• The FOREIGN KEY constraint is used to prevent actions that would destroy links between tables.</li> <li>• A FOREIGN KEY is a field in one table, that refers to the <u>PRIMARY KEY</u> in another table.</li> <li>• The following SQL creates a FOREIGN KEY on the "PersonID" column when the "Orders" table is created: <pre>CREATE TABLE Orders (     OrderID int NOT NULL,     OrderNumber int NOT NULL,     PersonID int,     PRIMARY KEY (OrderID),     FOREIGN KEY (PersonID) REFERENCES Persons(PersonID));</pre> </li> </ul> <p><b>SQL FOREIGN KEY on ALTER TABLE</b></p> <p>To create a FOREIGN KEY constraint on the "PersonID" column when the "Orders" table is already created, use the following SQL:</p> <pre>ALTER TABLE Orders ADD FOREIGN KEY (PersonID) REFERENCES Persons(PersonID);</pre> <p><b>DROP a FOREIGN KEY Constraint</b></p> <p>To drop a FOREIGN KEY constraint, use the following SQL:</p> <pre>ALTER TABLE Orders DROP FOREIGN KEY FK_PersonOrder;</pre> <p><b>iv) Default:-</b> Sets a default value for a column if no value is specified.</p> <ul style="list-style-type: none"> <li>• The <b>DEFAULT</b> keyword is used to set a default value for a column. When no value is specified for the column during an INSERT operation, the default value is automatically assigned.</li> </ul> <p><b>SQL DEFAULT on CREATE TABLE</b></p> <p>The following SQL sets a DEFAULT value for the "City" column when the "Persons" table is created:</p> <pre>CREATE TABLE Persons (     ID int NOT NULL,     LastName varchar(255) NOT NULL,     FirstName varchar(255),     Age int,     City varchar(255) DEFAULT 'Sandnes');</pre> <p><b>SQL DEFAULT on ALTER TABLE</b></p> <ul style="list-style-type: none"> <li>• To create a DEFAULT constraint on the "City" column when the table is already created, use the following SQL: <pre>ALTER TABLE Persons ALTER City SET DEFAULT 'Sandnes';</pre> </li> </ul> <p><b>DROP a DEFAULT Constraint</b></p> <ul style="list-style-type: none"> <li>• To drop a DEFAULT constraint, use the following SQL: <pre>ALTER TABLE Persons ALTER City DROP DEFAULT;</pre> </li> </ul> <p><b>v) Unique:-</b> Ensures that all values in a column are different.</p> <ul style="list-style-type: none"> <li>• The UNIQUE constraint ensures that all values in a column are different.</li> <li>• Both the UNIQUE and PRIMARY KEY constraints provide a guarantee for uniqueness for a column or set of columns.</li> <li>• A PRIMARY KEY constraint automatically has a UNIQUE constraint.</li> <li>• However, many UNIQUE constraints per table, but only one PRIMARY KEY constraint per table.</li> </ul> <p>The following SQL creates a UNIQUE constraint on the "ID" column when the "Persons" table is created:</p> <pre>CREATE TABLE Persons (     ID int NOT NULL,     LastName varchar(255) NOT NULL,     FirstName varchar(255),     Age int,     UNIQUE (ID));</pre> <p><b>SQL UNIQUE Constraint on ALTER TABLE</b></p> <ul style="list-style-type: none"> <li>• To create a UNIQUE constraint on the "ID" column when the table is already created, use the following SQL: <pre>ALTER TABLE Persons ADD UNIQUE (ID);</pre> </li> </ul> <p><b>DROP a UNIQUE Constraint</b></p> <p>To drop a UNIQUE constraint, use the following SQL:</p>	10	CO3	L2
---	---	----	-----	----



		ALTER TABLE Persons DROP INDEX UC_Person;			
6		<p>Consider the following tables: works (Pname, Cname, Salary) lives (Pname, Street, City) located-In (Cname, City) write the following queries in SQL:</p> <p>i) List the names of the people who work for the company 'Wipro' along with the cities they live in.</p> <p>ii) Find the names of the persons who do not work for 'Infosys'.</p> <p>iii) Find the people whose salaries are more than that of all of the 'oracle' employees.</p> <p>iv) Find the persons who works and lives in the same city. <b>(10 Marks)</b></p> <p>i) List the names of the people who work for the company 'Wipro' along with the cities they live in.</p> <p><code>SELECT w.Pname, l.City FROM works w JOIN lives l ON w.Pname = l.Pname WHERE w.Cname = 'Wipro';</code></p> <p>ii) Find the names of the persons who do not work for 'Infosys'.</p> <p><code>SELECT l.Pname FROM lives l WHERE l.Pname NOT IN ( SELECT w.Pname FROM works w WHERE w.Cname = 'Infosys' );</code></p> <p>iii) Find the people whose salaries are more than that of all of the 'Oracle' employees.</p> <p><code>SELECT w1.Pname FROM works w1 WHERE w1.Salary &gt; ALL ( SELECT w2.Salary FROM works w2 WHERE w2.Cname = 'Oracle' );</code></p> <p>iv) Find the persons who work and live in the same city.</p> <p><code>SELECT l.Pname FROM lives l JOIN located_In li ON l.City = li.City JOIN works w ON l.Pname = w.Pname AND w.Cname = li.Cname;</code></p>	10	CO3	L3

CI

CCI

HOD