

## Unit-5

### Introduction to transaction concepts

**Transaction** is a logical unit of work that represents real-world events of any organisation or an enterprise whereas concurrency control is the management of concurrent transaction execution. Transaction processing systems execute database transactions with large databases and hundreds of concurrent users, for example, railway and air reservations systems, banking system, credit card processing, stock market monitoring, super market inventory and checkouts and so on.

#### 1. TRANSACTION :

A **transaction** is a logical unit of work of database processing that includes one or more database access operations.

A transaction can be defined as an action or series of actions that is carried out by a single user or application program to perform operations for accessing the contents of the database. The operations can include retrieval, (Read), insertion (Write), deletion and modification. A transaction must be either completed or aborted.

It can either be embedded within an application program or can be specified interactively via a high-level query language such as SQL. Its execution preserves the consistency of the database. Each transaction should access shared data without interfering with the other transactions and whenever a transaction successfully completes its execution; its effect should be permanent. This basic abstraction frees the database application programmer from the following concerns :

- Inconsistencies caused by conflicting updates from concurrent users.
- Partially completed transactions in the event of systems failure.
- User-directed undoing of transactions.

A transaction is a sequence of READ and WRITE actions that are grouped together to form a database access. A transaction may consist of a simple SELECT operation to generate a list of table contents, or it may consist of a series of related UPDATE command sequences.

A transaction can include the following basic database access operations:

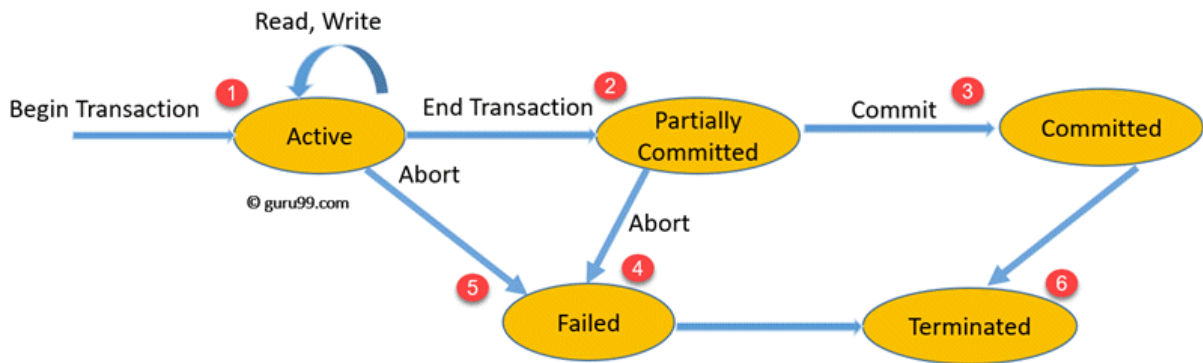
Operations	Descriptions
Retrieve	To retrieve data stored in a database.
Insert	To store new data in database.
Delete	To delete existing data from database.
Update	To modify existing data in database.
Commit	To save the work done permanently.

Rollback	To undo the work done.
----------	------------------------

Transaction that changes the contents of the database must alter the database from one **consistent** state to another. A consistent database state is one in which all data integrity constraints are satisfied. To ensure database consistency, every transaction must begin with the database in a known consistent state.

**Transaction Execution and Problems :**

A transaction which successfully completes its execution is said to have been committed. Otherwise, the transaction is aborted. Thus, if a committed transaction performs any update operation on the database, its effect must be reflected on the database even if there is a failure.



A transaction can be in one of the following states:

State	Description
Active state	A transaction goes into an active state immediately after it starts execution, where it can issue READ and WRITE operations. A transaction may be aborted when the transaction itself detects an error during execution which it cannot recover from, for example, a transaction trying to debit loan amount of an employee from his insufficient gross salary. A transaction may also be aborted before it has been committed due to system failure or any other circumstances beyond its control.
Partially committed	When the transaction ends, it moves to the partially committed state. When the last state is reached. To this point, some recovery protocols need to ensure that a system failure will not result in an inability to record the changes of the transaction permanently. Once this check is successful, the transaction is said to have reached its commit point and enters the committed state.
Aborted	When the normal execution can no longer be performed. Failed or aborted transactions may be restarted later, either automatically or after being resubmitted by the user as new transactions.
Committed	After successful completion of transaction.

A transaction is said to be in a committed state if it has partially committed and it can be ensured that it will never be aborted.
---

### **Transaction Execution with SQL :**

The American National Standards Institute (ANSI) has defined standards that govern SQL database transactions. Transaction support is provided by two SQL statements namely COMMIT and ROLLBACK.

The ANSI standards require that, when a transaction sequence is initiated by a user or an application program, it must continue through all succeeding SQL statements until one of the following four events occur :

- A COMMIT statement is reached, in which case all changes are permanently recorded within the database. The COMMIT statement automatically ends the SQL transaction. The COMMIT operations indicates successful end-of-transaction.
- A ROLLBACK statement is reached, in which case all the changes are aborted and the database is rolled back to its previous consistent state. The ROLLBACK operation indicates unsuccessful end-of-transaction.
- The end of a program is successfully reached, in which case all changes are permanently recorded within the database. This action is equivalent to COMMIT.
- The program is abnormally terminated, in which case the changes made in the database are aborted and the database is rolled back to its previous consistent state. This action is equivalent to ROLLBACK.

### **Transaction Properties :**

A transaction must have the following four properties, called ACID properties (also called ACIDITY of a transaction), to ensure that a database remains stable state after the transaction is executed:

1. Atomicity.
2. Consistency.
3. Isolation.
4. Durability.

#### **Atomicity:**

The atomicity property of a transaction requires that all operations of a transaction be completed, if not, the transaction is aborted. In other words, a transaction is treated as single, individual logical unit of work.

Therefore, a transaction must execute and complete each operation in its logic before it commits its changes. As stated earlier, the transaction is considered as one operation even though there are multiple read and writes. Thus, transaction completes or fails as one unit. The atomicity property of transaction is ensured by the transaction recovery subsystem of a DBMS. In the event of a system crash in the midst of transaction execution, the recovery techniques undo any effects of the transaction on the database.

#### **Consistency:**

Database consistency is the property that every transaction sees a consistent database instance. In other words, execution of a transaction must leave a database in either its prior stable state or a new stable state that reflects the new modifications (updates) made by the transaction. If the transaction fails, the database must be returned to the state it was in prior to the execution of the

failed transaction.  
If the transaction commits, the database must reflect the new changes. Thus, all resources are always in a consistent state. The preservation of consistency is generally the responsibility of the programmers who write the database programs or of the DBMS module that enforces integrity constraints. A database program should be written in a way that guarantees that, if the database is in a consistent state before executing the transaction, it will be in a consistent state after the complete execution of the transaction, assuming that no interference with other transactions occur. In other words, a transaction must transform the database from one consistent state to another consistent state.

### **Isolation:**

Isolation property of a transaction means that the data used during the execution of a transaction cannot be used by a second transaction until the first one is completed. This property isolates transactions from one another. In other words, if a transaction T1 is being executed and is using the data item X, that data item cannot be accessed by any other transaction (T2.....Tn) until T1 ends. The isolation property is enforced by the concurrency control subsystem of the DBMS.

### **Durability:**

The durability property of transaction indicates the performance of the database's consistent state. It states that the changes made by a transaction are permanent. They cannot be lost by either a system failure or by the erroneous operation of a faulty transaction. When a transaction is completed, the database reaches a consistent state and that state cannot be lost, even in the event of system's failure. Durability property is the responsibility of the recovery subsystem of the DBMS.

## **2.IMPLEMENTATION OF ATOMICITY AND DURABILITY:**


The recovery-management component of a database system can support atomicity and durability by a variety of schemes.

- Here we are going to learn about one of the simplest scheme called Shadow copy. In the shadow-copy scheme, a transaction that wants to update the database first creates a complete copy of the database. All updates are done on the new database copy, leaving the original copy, the shadow copy, untouched. If at any point the transaction has to be aborted, the system merely deletes the new copy. The old copy of the database has not been affected.
- This scheme is based on making copies of the database, called shadow copies, assumes that only one transaction is active at a time. The scheme also assumes that the database is simply a file on disk. A pointer called db-pointer is maintained on disk; it points to the current copy of the database.

If the transaction completes, it is committed as follows:

- First, the operating system is asked to make sure that all pages of the new copy of the database have been written out to disk. (Unix systems use the flush command for this purpose.)
- After the operating system has written all the pages to disk, the database system updates the pointer db-pointer to point to the new copy of the database; the new copy then becomes the current copy of the database. The old copy of the database is then deleted.

Figure below depicts the scheme, showing the database state before and after the update.



The transaction is said to have been committed at the point where the updated db pointer is written to disk.

#### **How the technique handles transaction failures:**

- If the transaction fails at any time before db-pointer is updated, the old contents of the database are not affected.
- We can abort the transaction by just deleting the new copy of the database.
- Once the transaction has been committed, all the updates that it performed are in the database pointed to by db pointer.
- Thus, either all updates of the transaction are reflected, or none of the effects are reflected, regardless of transaction failure.

#### **How the technique handles system failures:**

- Suppose that the system fails at any time before the updated db-pointer is written to disk. Then, when the system restarts, it will read db-pointer and will thus see the original contents of the database, and none of the effects of the transaction will be visible on the database.
- Next, suppose that the system fails after db-pointer has been updated on disk. Before the pointer is updated, all updated pages of the new copy of the database were written to disk.
- Again, we assume that, once a file is written to disk, its contents will not be damaged even if there is a system failure. Therefore, when the system restarts, it will read db-pointer and will thus see the contents of the database after all the updates performed by the transaction.

### **3.CONCURRENT EXECUTION IN TRANSACTION:**

Transaction-processing systems usually allow multiple transactions to run concurrently. Allowing multiple transactions to update data concurrently causes several complications with consistency of the data.

Ensuring consistency in spite of concurrent execution of transactions requires extra work; it is far easier to insist that transactions run serially—that is, one at a time, each starting only after the previous one has completed.

However, there are two good reasons for allowing concurrency:

Improved throughput and resource utilization:

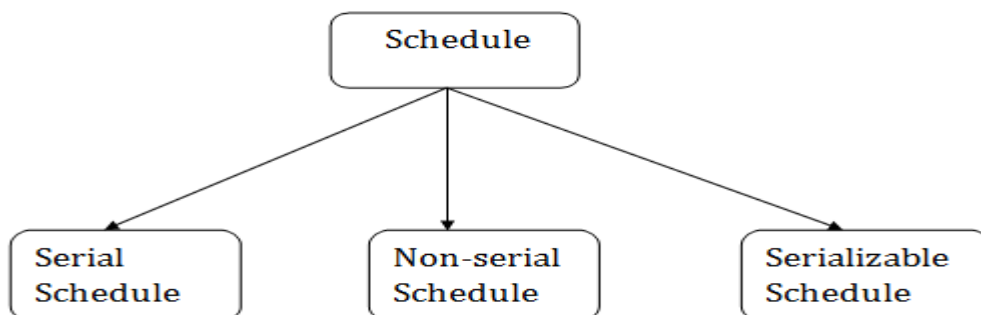
- A transaction consists of many steps. Some involve I/O activity; others involve CPU activity. The CPU and the disks in a computer system can operate in parallel. Therefore, I/O activity can be done in parallel with processing at the CPU.
- The parallelism of the CPU and the I/O system can therefore be exploited to run multiple transactions in parallel.
- While a read or write on behalf of one transaction is in progress on one disk, another transaction can be running in the CPU, while another disk may be executing a read or write on behalf of a third transaction.
- All of this increases the throughput of the system—that is, the number of transactions executed in a given amount of time.
- Correspondingly, the processor and disk utilization also increase; in other words, the processor and disk spend less time idle, or not performing any useful work.

Reduced waiting time:

- There may be a mix of transactions running on a system, some short and some long.
- If transactions run serially, a short transaction may have to wait for a preceding long transaction to complete, which can lead to unpredictable delays in running a transaction.
- If the transactions are operating on different parts of the database, it is better to let them run concurrently, sharing the CPU cycles and disk accesses among them.
- Concurrent execution reduces the unpredictable delays in running transactions.
- Moreover, it also reduces the average response time: the average time for a transaction to be completed after it has been submitted.

#### 4.Schedule

A series of operation from one transaction to another transaction is known as schedule. It is used to preserve the order of the operation in each of the individual transaction.



## 1. Serial Schedule

The serial schedule is a type of schedule where one transaction is executed completely before starting another transaction. In the serial schedule, when the first transaction completes its cycle, then the next transaction is executed.

**For example:** Suppose there are two transactions T1 and T2 which have some operations. If it has no interleaving of operations, then there are the following two possible outcomes:

1. Execute all the operations of T1 which was followed by all the operations of T2.
  2. Execute all the operations of T2 which was followed by all the operations of T1.
- In the given (a) figure, Schedule A shows the serial schedule where T1 followed by T2.
  - In the given (b) figure, Schedule B shows the serial schedule where T2 followed by T1.


## 2. Non-serial Schedule

- If interleaving of operations is allowed, then there will be non-serial schedule.
- It contains many possible orders in which the system can execute the individual operations of the transactions.
- In the given figure (c) and (d), Schedule C and Schedule D are the non-serial schedules. It has interleaving of operations.

## 3. Serializable schedule

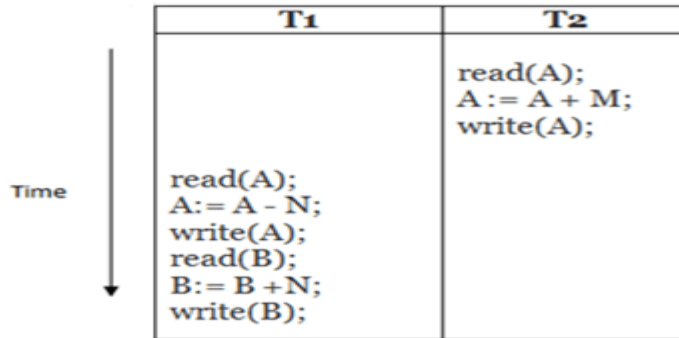
- The serializability of schedules is used to find non-serial schedules that allow the transaction to execute concurrently without interfering with one another.
- It identifies which schedules are correct when executions of the transaction have interleaving of their operations.
- A non-serial schedule will be serializable if its result is equal to the result of its transactions executed serially.

(a)



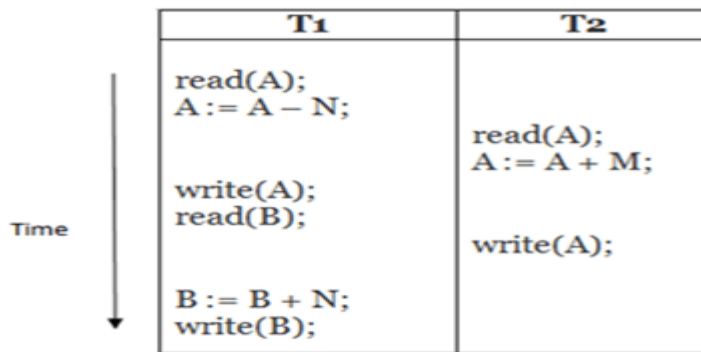
Schedule A

(b)




**Schedule B**

(c)



**Schedule C**

(d)



**Schedule D**


## 5. Serializability

Serialization Graph is used to test the Serializability of a schedule.

Assume a schedule S. For S, we construct a graph known as precedence graph. This graph has a pair  $G = (V, E)$ , where V consists a set of vertices, and E consists a set of edges. The set of vertices is used to contain all the transactions participating in the schedule. The set of edges is used to contain all edges  $T_i \rightarrow T_j$  for which one of the three conditions holds:

1. Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes write (Q) before  $T_j$  executes read (Q).
2. Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes read (Q) before  $T_j$  executes write (Q).
3. Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes write (Q) before  $T_j$  executes write (Q).

### Precedence graph for Schedule S



- If a precedence graph contains a single edge  $T_i \rightarrow T_j$ , then all the instructions of  $T_i$  are executed before the first instruction of  $T_j$  is executed.
- If a precedence graph for schedule S contains a cycle, then S is non-serializable. If the precedence graph has no cycle, then S is known as serializable.

For example:


	<b>T1</b>	<b>T2</b>	<b>T3</b>
<b>Time</b> 	Read(A)	Read(B)	Read(C)
	A := f <sub>1</sub> (A)	B := f <sub>2</sub> (B) Write(B)	C := f <sub>3</sub> (C) Write(C)
	Write(A)	Read(A) A := f <sub>4</sub> (A)	Read(B)
	Read(C)	Write(A)	
	C := f <sub>5</sub> (C) Write(C)		B := f <sub>6</sub> (B) Write(B)

**Schedule S1**

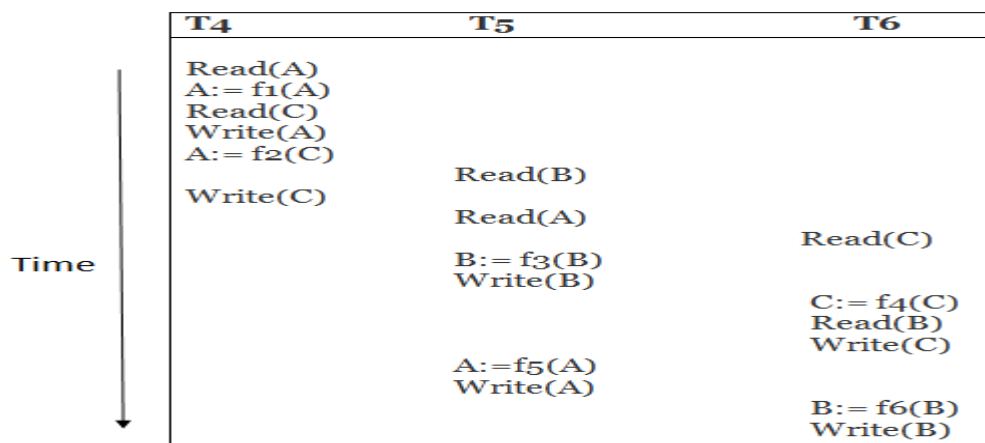
**Explanation:**

**Read(A):** In T1, no subsequent writes to A, so no new edges  
**Read(B):** In T2, no subsequent writes to B, so no new edges  
**Read(C):** In T3, no subsequent writes to C, so no new edges  
**Write(B):** B is subsequently read by T3, so add edge T2 → T3  
**Write(C):** C is subsequently read by T1, so add edge T3 → T1  
**Write(A):** A is subsequently read by T2, so add edge T1 → T2  
**Write(A):** In T2, no subsequent reads to A, so no new edges  
**Write(C):** In T1, no subsequent reads to C, so no new edges  
**Write(B):** In T3, no subsequent reads to B, so no new edges

Precedence graph for schedule S1:



The precedence graph for schedule S1 contains a cycle that's why Schedule S1 is non-serializable.




**Schedule S2**

**Explanation:**

**Read(A):** In T4, no subsequent writes to A, so no new edges  
**Read(C):** In T4, no subsequent writes to C, so no new edges  
**Write(A):** A is subsequently read by T5, so add edge T4 → T5  
**Read(B):** In T5, no subsequent writes to B, so no new edges  
**Write(C):** C is subsequently read by T6, so add edge T4 → T6  
**Write(B):** A is subsequently read by T6, so add edge T5 → T6  
**Write(C):** In T6, no subsequent reads to C, so no new edges  
**Write(A):** In T5, no subsequent reads to A, so no new edges  
**Write(B):** In T6, no subsequent reads to B, so no new edges

Precedence graph for schedule S2:



The precedence graph for schedule S2 contains no cycle that's why Schedule S2 is serializable.

### Conflict Serializable Schedule

- A schedule is called conflict serializable if after swapping of non-conflicting operations, it can transform into a serial schedule.
- The schedule will be a conflict serializable if it is conflict equivalent to a serial schedule.

### Conflicting Operations

The two operations become conflicting if all conditions satisfy:

1. Both belong to separate transactions.
2. They have the same data item.
3. They contain at least one write operation.

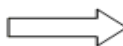
Example:

Swapping is possible only if S1 and S2 are logically equal.

#### 1. T1: Read(A) T2: Read(A)

T1	T2
Read(A)	
	Read(A)

**Schedule S1**


Swapped  


T1	T2
	Read(A)
Read(A)	

**Schedule S2**

Here,  $S1 = S2$ . That means it is non-conflict.

## 2. T1: Read(A) T2: Write(A)



**Schedule S1**

**Schedule S2**

Here,  $S1 \neq S2$ . That means it is conflict.

### Conflict Equivalent

In the conflict equivalent, one can be transformed to another by swapping non-conflicting operations. In the given example, S2 is conflict equivalent to S1 (S1 can be converted to S2 by swapping non-conflicting operations).

Two schedules are said to be conflict equivalent if and only if:

1. They contain the same set of the transaction.
2. If each pair of conflict operations are ordered in the same way.

Example:

### Non-serial schedule

T1	T2
Read(A) Write(A)	Read(A) Write(A)
Read(B) Write(B)	

**Schedule S1**

### Serial Schedule

T1	T2
Read(A) Write(A) Read(B) Write(B)	Read(A) Write(A) Read(B) Write(B)

**Schedule S2**

Schedule S2 is a serial schedule because, in this, all operations of T1 are performed before starting any operation of T2. Schedule S1 can be transformed into a serial schedule by swapping non-conflicting operations of S1.

**After swapping of non-conflict operations, the schedule S1 becomes:**

T1	T2
Read(A) Write(A) Read(B) Write(B)	Read(A) Write(A) Read(B) Write(B)

Since, S1 is conflict serializable.

### 6. View Serializability

- A schedule will view serializable if it is view equivalent to a serial schedule.
- If a schedule is conflict serializable, then it will be view serializable.
- The view serializable which does not conflict serializable contains blind writes.

#### View Equivalent

Two schedules S1 and S2 are said to be view equivalent if they satisfy the following conditions:

##### 1. Initial Read

An initial read of both schedules must be the same. Suppose two schedule S1 and S2. In schedule S1, if a transaction T1 is reading the data item A, then in S2, transaction T1 should also read A.

T1	T2
Read(A)	Write(A)

**Schedule S1**

T1	T2
Read(A)	Write(A)

**Schedule S2**

Above two schedules are view equivalent because Initial read operation in S1 is done by T1 and in S2 it is also done by T1.

## 2. Updated Read

In schedule S1, if  $T_i$  is reading A which is updated by  $T_j$  then in S2 also,  $T_i$  should read A which is updated by  $T_j$ .

T1	T2	T3
Write(A)	Write(A)	Read(A)

**Schedule S1**

T1	T2	T3
Write(A)	Write(A)	<u>Read(A)</u>

**Schedule S2**

Above two schedules are not view equal because, in S1, T3 is reading A updated by T2 and in S2, T3 is reading A updated by T1.

## 3. Final Write

A final write must be the same between both the schedules. In schedule S1, if a transaction T1 updates A at last then in S2, final writes operations should also be done by T1.

T1	T2	T3
Write(A)	Read(A)	Write(A)

**Schedule S1**

T1	T2	T3
Write(A)	Read(A)	Write(A)

**Schedule S2**

Above two schedules is view equal because Final write operation in S1 is done by T3 and in S2, the final write operation is also done by T3.

### Example:

T1	T2	T3
Read(A)	Write(A)	
Write(A)		Write(A)

## Schedule S

With 3 transactions, the total number of possible schedule

1.  $= 3! = 6$
2. S1 = <T1 T2 T3>
3. S2 = <T1 T3 T2>
4. S3 = <T2 T3 T1>
5. S4 = <T2 T1 T3>
6. S5 = <T3 T1 T2>
7. S6 = <T3 T2 T1>

**Taking first schedule S1:**

T1	T2	T3
Read(A) Write(A)	Write(A)	Write(A)

## Schedule S1

**Step 1:** final updation on data items

In both schedules S and S1, there is no read except the initial read that's why we don't need to check that condition.

**Step 2:** Initial Read

The initial read operation in S is done by T1 and in S1, it is also done by T1.

**Step 3:** Final Write

The final write operation in S is done by T3 and in S1, it is also done by T3. So, S and S1 are view Equivalent.

The first schedule S1 satisfies all three conditions, so we don't need to check another schedule.

1. T1 → T2 → T3

## Recoverability of Schedule

Sometimes a transaction may not execute completely due to a software issue, system crash or hardware failure. In that case, the failed transaction has to be rollback. But some other transaction

may also have used value produced by the failed transaction. So we also have to rollback those transactions.

T1	T1's buffer space	T2	T2's buffer space	Database
				A = 6500
Read(A);	A = 6500			A = 6500
A = A - 500;	A = 6000			A = 6500
Write(A);	A = 6000			A = 6000
		Read(A);	A = 6000	A = 6000
		A = A + 1000;	A = 7000	A = 6000
		Write(A);	A = 7000	A = 7000
		Commit;		
Failure Point				
Commit;				

The above table 1 shows a schedule which has two transactions. T1 reads and writes the value of A and that value is read and written by T2. T2 commits but later on, T1 fails. Due to the failure, we have to rollback T1. T2 should also be rollback because it reads the value written by T1, but T2 can't be rollback because it already committed. So this type of schedule is known as irrecoverable schedule.

**Irrecoverable schedule:** The schedule will be irrecoverable if Tj reads the updated value of Ti and Tj committed before Ti commit.

T1	T1's buffer space	T2	T2's buffer space	Database
				A = 6500
Read(A);	A = 6500			A = 6500
A = A - 500;	A = 6000			A = 6500
Write(A);	A = 6000			A = 6000
		Read(A);	A = 6000	A = 6000
		A = A + 1000;	A = 7000	A = 6000
		Write(A);	A = 7000	A = 7000
Failure Point				
Commit;				
		Commit;		

The above table 2 shows a schedule with two transactions. Transaction T1 reads and writes A, and that value is read and written by transaction T2. But later on, T1 fails. Due to this, we have to rollback T1. T2 should be rollback because T2 has read the value written by T1. As it has not committed before T1 commits so we can rollback transaction T2 as well. So it is recoverable with cascade rollback.

**Recoverable with cascading rollback:** The schedule will be recoverable with cascading rollback if Tj reads the updated value of Ti. Commit of Tj is delayed till commit of Ti.

T1	T1's buffer space	T2	T2's buffer space	Database
				A = 6500
Read(A);	A = 6500			A = 6500
A = A - 500;	A = 6000			A = 6500
Write(A);	A = 6000			A = 6000
Commit;		Read(A);	A = 6000	A = 6000
		A = A + 1000;	A = 7000	A = 6000
		Write(A);	A = 7000	A = 7000
		Commit;		

The above Table 3 shows a schedule with two transactions. Transaction T1 reads and write A and commits, and that value is read and written by T2. So this is a cascade less recoverable schedule.

### Failure Classification

To find that where the problem has occurred, we generalize a failure into the following categories:

1. Transaction failure
2. System crash
3. Disk failure

#### 1. Transaction failure

The transaction failure occurs when it fails to execute or when it reaches a point from where it can't go any further. If a few transaction or process is hurt, then this is called as transaction failure.

Reasons for a transaction failure could be -

1. **Logical errors:** If a transaction cannot complete due to some code error or an internal error condition, then the logical error occurs.
2. **Syntax error:** It occurs where the DBMS itself terminates an active transaction because the database system is not able to execute it. **For example,** The system aborts an active transaction, in case of deadlock or resource unavailability.

#### 2. System Crash

- System failure can occur due to power failure or other hardware or software failure. **Example:** Operating system error.

**Fail-stop assumption:** In the system crash, non-volatile storage is assumed not to be corrupted.

### 3. Disk Failure

- It occurs where hard-disk drives or storage drives used to fail frequently. It was a common problem in the early days of technology evolution.
- Disk failure occurs due to the formation of bad sectors, disk head crash, and unreachability to the disk or any other failure, which destroy all or part of disk storage.

### Log-Based Recovery

- The log is a sequence of records. Log of each transaction is maintained in some stable storage so that if any failure occurs, then it can be recovered from there.
- If any operation is performed on the database, then it will be recorded in the log.
- But the process of storing the logs should be done before the actual transaction is applied in the database.

Let's assume there is a transaction to modify the City of a student. The following logs are written for this transaction.

- When the transaction is initiated, then it writes 'start' log.
  1. <Tn, Start>
- When the transaction modifies the City from 'Noida' to 'Bangalore', then another log is written to the file.
  1. <Tn, City, 'Noida', 'Bangalore' >
- When the transaction is finished, then it writes another log to indicate the end of the transaction.
  1. <Tn, Commit>

There are two approaches to modify the database:

#### 1. Deferred database modification:

- The deferred modification technique occurs if the transaction does not modify the database until it has committed.
- In this method, all the logs are created and stored in the stable storage, and the database is updated when a transaction commits.

## 2. Immediate database modification:

- The Immediate modification technique occurs if database modification occurs while the transaction is still active.
- In this technique, the database is modified immediately after every operation. It follows an actual database modification.

## Recovery using Log records

When the system is crashed, then the system consults the log to find which transactions need to be undone and which need to be redone.


1. If the log contains the record  $\langle T_i, \text{Start} \rangle$  and  $\langle T_i, \text{Commit} \rangle$  or  $\langle T_i, \text{Commit} \rangle$ , then the Transaction  $T_i$  needs to be redone.
2. If log contains record  $\langle T_n, \text{Start} \rangle$  but does not contain the record either  $\langle T_i, \text{commit} \rangle$  or  $\langle T_i, \text{abort} \rangle$ , then the Transaction  $T_i$  needs to be undone.

## 7. Deadlock in DBMS

A deadlock is a condition where two or more transactions are waiting indefinitely for one another to give up locks. Deadlock is said to be one of the most feared complications in DBMS as no task ever gets finished and is in waiting state forever.

**For example:** In the student table, transaction  $T_1$  holds a lock on some rows and needs to update some rows in the grade table. Simultaneously, transaction  $T_2$  holds locks on some rows in the grade table and needs to update the rows in the Student table held by Transaction  $T_1$ .

Now, the main problem arises. Now Transaction  $T_1$  is waiting for  $T_2$  to release its lock and similarly, transaction  $T_2$  is waiting for  $T_1$  to release its lock. All activities come to a halt state and remain at a standstill. It will remain in a standstill until the DBMS detects the deadlock and aborts one of the transactions.



**Figure: Deadlock in DBMS**

## Deadlock Avoidance

- When a database is stuck in a deadlock state, then it is better to avoid the database rather than aborting or restating the database. This is a waste of time and resource.

- Deadlock avoidance mechanism is used to detect any deadlock situation in advance. A method like "wait for graph" is used for detecting the deadlock situation but this method is suitable only for the smaller database. For the larger database, deadlock prevention method can be used.


### Deadlock Detection

In a database, when a transaction waits indefinitely to obtain a lock, then the DBMS should detect whether the transaction is involved in a deadlock or not. The lock manager maintains a Wait for the graph to detect the deadlock cycle in the database.

### Wait for Graph

- This is the suitable method for deadlock detection. In this method, a graph is created based on the transaction and their lock. If the created graph has a cycle or closed loop, then there is a deadlock.
- The wait for the graph is maintained by the system for every transaction which is waiting for some data held by the others. The system keeps checking the graph if there is any cycle in the graph.

The wait for a graph for the above scenario is shown below:



### Deadlock Prevention

- Deadlock prevention method is suitable for a large database. If the resources are allocated in such a way that deadlock never occurs, then the deadlock can be prevented.

- The Database management system analyzes the operations of the transaction whether they can create a deadlock situation or not. If they do, then the DBMS never allowed that transaction to be executed.

### Wait-Die scheme

In this scheme, if a transaction requests for a resource which is already held with a conflicting lock by another transaction then the DBMS simply checks the timestamp of both transactions. It allows the older transaction to wait until the resource is available for execution.

Let's assume there are two transactions  $T_i$  and  $T_j$  and let  $TS(T)$  is a timestamp of any transaction  $T$ . If  $T_2$  holds a lock by some other transaction and  $T_1$  is requesting for resources held by  $T_2$  then the following actions are performed by DBMS:

1. Check if  $TS(T_i) < TS(T_j)$  - If  $T_i$  is the older transaction and  $T_j$  has held some resource, then  $T_i$  is allowed to wait until the data-item is available for execution. That means if the older transaction is waiting for a resource which is locked by the younger transaction, then the older transaction is allowed to wait for resource until it is available.
2. Check if  $TS(T_i) < TS(T_j)$  - If  $T_i$  is older transaction and has held some resource and if  $T_j$  is waiting for it, then  $T_j$  is killed and restarted later with the random delay but with the same timestamp.

### Wound wait scheme

- In wound wait scheme, if the older transaction requests for a resource which is held by the younger transaction, then older transaction forces younger one to kill the transaction and release the resource. After the minute delay, the younger transaction is restarted but with the same timestamp.
- If the older transaction has held a resource which is requested by the Younger transaction, then the younger transaction is asked to wait until older releases it.

## 8.Concurrency Control

Concurrency Control is the management procedure that is required for controlling concurrent execution of the operations that take place on a database.

But before knowing about concurrency control, we should know about concurrent execution.

### Concurrent Execution in DBMS

- In a multi-user system, multiple users can access and use the same database at one time, which is known as the concurrent execution of the database. It means that the same database is executed simultaneously on a multi-user system by different users.

- While working on the database transactions, there occurs the requirement of using the database by multiple users for performing different operations, and in that case, concurrent execution of the database is performed.
- The thing is that the simultaneous execution that is performed should be done in an interleaved manner, and no operation should affect the other executing operations, thus maintaining the consistency of the database. Thus, on making the concurrent execution of the transaction operations, there occur several challenging problems that need to be solved.

### Problems with Concurrent Execution

In a database transaction, the two main operations are **READ** and **WRITE** operations. So, there is a need to manage these two operations in the concurrent execution of the transactions as if these operations are not performed in an interleaved manner, and the data may become inconsistent. So, the following problems occur with the Concurrent Execution of the operations:

#### Problem 1: Lost Update Problems (W - W Conflict)

The problem occurs *when two different database transactions perform the read/write operations on the same database items in an interleaved manner (i.e., concurrent execution) that makes the values of the items incorrect hence making the database inconsistent.*

**For example:**

Consider the below diagram where two transactions  $T_X$  and  $T_Y$ , are performed on the same account A where the balance of account A is \$300.

Time	$T_X$	$T_Y$
$t_1$	READ (A)	—
$t_2$	$A = A - 50$	—
$t_3$	—	READ (A)
$t_4$	—	$A = A + 100$
$t_5$	—	—
$t_6$	WRITE (A)	—
$t_7$	—	WRITE (A)

**LOST UPDATE PROBLEM**

- At time  $t_1$ , transaction  $T_X$  reads the value of account A, i.e., \$300 (only read).
- At time  $t_2$ , transaction  $T_X$  deducts \$50 from account A that becomes \$250 (only deducted and not updated/write).
- Alternately, at time  $t_3$ , transaction  $T_Y$  reads the value of account A that will be \$300 only because  $T_X$  didn't update the value yet.
- At time  $t_4$ , transaction  $T_Y$  adds \$100 to account A that becomes \$400 (only added but not updated/write).
- At time  $t_6$ , transaction  $T_X$  writes the value of account A that will be updated as \$250 only, as  $T_Y$  didn't update the value yet.
- Similarly, at time  $t_7$ , transaction  $T_Y$  writes the values of account A, so it will write as done at time  $t_4$  that will be \$400. It means the value written by  $T_X$  is lost, i.e., \$250 is lost.

Hence data becomes incorrect, and database sets to inconsistent.

### Dirty Read Problems (W-R Conflict)

The dirty read problem occurs *when one transaction updates an item of the database, and somehow the transaction fails, and before the data gets rollback, the updated database item is accessed by another transaction. There comes the Read-Write Conflict between both transactions.*

**For example:**

**Consider two transactions  $T_X$  and  $T_Y$  in the below diagram performing read/write operations on account A where the available balance in account A is \$300:**

Time	$T_X$	$T_Y$
$t_1$	READ (A)	—
$t_2$	$A = A + 50$	—
$t_3$	WRITE (A)	—
$t_4$	—	READ (A)
$t_5$	SERVER DOWN ROLLBACK	—

### DIRTY READ PROBLEM

- At time  $t_1$ , transaction  $T_X$  reads the value of account A, i.e., \$300.
- At time  $t_2$ , transaction  $T_X$  adds \$50 to account A that becomes \$350.
- At time  $t_3$ , transaction  $T_X$  writes the updated value in account A, i.e., \$350.
- Then at time  $t_4$ , transaction  $T_Y$  reads account A that will be read as \$350.
- Then at time  $t_5$ , transaction  $T_X$  rollbacks due to server problem, and the value changes back to \$300 (as initially).
- But the value for account A remains \$350 for transaction  $T_Y$  as committed, which is the dirty read and therefore known as the Dirty Read Problem.

### Unrepeatable Read Problem (W-R Conflict)

Also known as *Inconsistent Retrievals Problem* that occurs when in a transaction, two different values are read for the same database item.

Consider two transactions,  $T_X$  and  $T_Y$ , performing the read/write operations on account A, having an available balance = \$300. The diagram is shown below:

Time	$T_X$	$T_Y$
$t_1$	READ (A)	—
$t_2$	—	READ (A)
$t_3$	—	$A = A + 100$
$t_4$	—	WRITE (A)
$t_5$	READ (A)	—

### UNREPEATABLE READ PROBLEM

- At time  $t_1$ , transaction  $T_X$  reads the value from account A, i.e., \$300.
- At time  $t_2$ , transaction  $T_Y$  reads the value from account A, i.e., \$300.
- At time  $t_3$ , transaction  $T_Y$  updates the value of account A by adding \$100 to the available balance, and then it becomes \$400.
- At time  $t_4$ , transaction  $T_Y$  writes the updated value, i.e., \$400.

- After that, at time  $t_5$ , transaction  $T_X$  reads the available value of account A, and that will be read as \$400.
- It means that within the same transaction  $T_X$ , it reads two different values of account A, i.e., \$300 initially, and after updation made by transaction  $T_Y$ , it reads \$400. It is an unrepeatable read and is therefore known as the Unrepeatable read problem.

Thus, in order to maintain consistency in the database and avoid such problems that take place in concurrent execution, management is needed, and that is where the concept of Concurrency Control comes into role.

## 9. Concurrency Control

Concurrency Control is the working concept that is required for controlling and managing the concurrent execution of database operations and thus avoiding the inconsistencies in the database. Thus, for maintaining the concurrency of the database, we have the concurrency control protocols.

### Concurrency Control Protocols

The concurrency control protocols ensure the *atomicity*, *consistency*, *isolation*, *durability* and *serializability* of the concurrent execution of the database transactions. Therefore, these protocols are categorized as:

- Lock Based Concurrency Control Protocol
- Time Stamp Concurrency Control Protocol

### Validation Based Concurrency Control Protocol Lock-Based Protocol

In this type of protocol, any transaction cannot read or write data until it acquires an appropriate lock on it. There are two types of lock:

#### 1. Shared lock:

- It is also known as a Read-only lock. In a shared lock, the data item can only read by the transaction.
- It can be shared between the transactions because when the transaction holds a lock, then it can't update the data on the data item.

#### 2. Exclusive lock:

- In the exclusive lock, the data item can be both reads as well as written by the transaction.
- This lock is exclusive, and in this lock, multiple transactions do not modify the same data simultaneously.


There are four types of lock protocols available:

### 1. Simplistic lock protocol

It is the simplest way of locking the data while transaction. Simplistic lock-based protocols allow all the transactions to get the lock on the data before insert or delete or update on it. It will unlock the data item after completing the transaction.


### 2. Pre-claiming Lock Protocol

- Pre-claiming Lock Protocols evaluate the transaction to list all the data items on which they need locks.
- Before initiating an execution of the transaction, it requests DBMS for all the lock on all those data items.
- If all the locks are granted then this protocol allows the transaction to begin. When the transaction is completed then it releases all the lock.
- If all the locks are not granted then this protocol allows the transaction to rolls back and waits until all the locks are granted.



### 3. Two-phase locking (2PL)

- The two-phase locking protocol divides the execution phase of the transaction into three parts.
- In the first part, when the execution of the transaction starts, it seeks permission for the lock it requires.
- In the second part, the transaction acquires all the locks. The third phase is started as soon as the transaction releases its first lock.
- In the third phase, the transaction cannot demand any new locks. It only releases the acquired locks.



There are two phases of 2PL:

**Growing phase:** In the growing phase, a new lock on the data item may be acquired by the transaction, but none can be released.

**Shrinking phase:** In the shrinking phase, existing lock held by the transaction may be released, but no new locks can be acquired.

In the below example, if lock conversion is allowed then the following phase can happen:

1. Upgrading of lock (from S(a) to X (a)) is allowed in growing phase.
2. Downgrading of lock (from X(a) to S(a)) must be done in shrinking phase.

**Example:**

	<b>T1</b>	<b>T2</b>
0	LOCK-S(A)	
1		LOCK-S(A)
2	LOCK-X(B)	
3	---	---
4	UNLOCK(A)	
5		LOCK-X(C)
6	UNLOCK(B)	
7		UNLOCK(A)
8		UNLOCK(C)
9	---	---

The following way shows how unlocking and locking work with 2-PL.

**Transaction T1:**


- **Growing phase:** from step 1-3
- **Shrinking phase:** from step 5-7
- **Lock point:** at 3

**Transaction T2:**

- **Growing phase:** from step 2-6
- **Shrinking phase:** from step 8-9
- **Lock point:** at 6

4. Strict Two-phase locking (Strict-2PL)


- The first phase of Strict-2PL is similar to 2PL. In the first phase, after acquiring all the locks, the transaction continues to execute normally.
- The only difference between 2PL and strict 2PL is that Strict-2PL does not release a lock after using it.
- Strict-2PL waits until the whole transaction to commit, and then it releases all the locks at a time.
- Strict-2PL protocol does not have shrinking phase of lock release.



- **Growing phase:** from step 2-6
- **Shrinking phase:** from step 8-9
- **Lock point:** at 6

#### 4. Strict Two-phase locking (Strict-2PL)

- The first phase of Strict-2PL is similar to 2PL. In the first phase, after acquiring all the locks, the transaction continues to execute normally.
- The only difference between 2PL and strict 2PL is that Strict-2PL does not release a lock after using it.
- Strict-2PL waits until the whole transaction to commit, and then it releases all the locks at a time.
- Strict-2PL protocol does not have shrinking phase of lock release.



$TS(T_i)$  denotes the timestamp of the transaction  $T_i$ .

$R\_TS(X)$  denotes the Read time-stamp of data-item  $X$ .

$W\_TS(X)$  denotes the Write time-stamp of data-item  $X$ .

#### Validation Based Protocol

Validation phase is also known as optimistic concurrency control technique. In the validation based protocol, the transaction is executed in the following three phases:

1. **Read phase:** In this phase, the transaction  $T$  is read and executed. It is used to read the value of various data items and stores them in temporary local variables. It can perform all the write operations on temporary variables without an update to the actual database.
2. **Validation phase:** In this phase, the temporary variable value will be validated against the actual data to see if it violates the serializability.
3. **Write phase:** If the validation of the transaction is validated, then the temporary results are written to the database or system otherwise the transaction is rolled back.

Here each phase has the following different timestamps:

**Start( $T_i$ ):** It contains the time when  $T_i$  started its execution.

**Validation ( $T_i$ ):** It contains the time when  $T_i$  finishes its read phase and starts its validation phase.

**Finish( $T_i$ ):** It contains the time when  $T_i$  finishes its write phase.

- This protocol is used to determine the time stamp for the transaction for serialization using the time stamp of the validation phase, as it is the actual phase which determines if the transaction will commit or rollback.
- Hence  $TS(T) = \text{validation}(T)$ .
- The serializability is determined during the validation process. It can't be decided in advance.
- While executing the transaction, it ensures a greater degree of concurrency and also less number of conflicts.
- Thus it contains transactions which have less number of rollbacks.

**9.Granularity:** It is the size of data item allowed to lock.

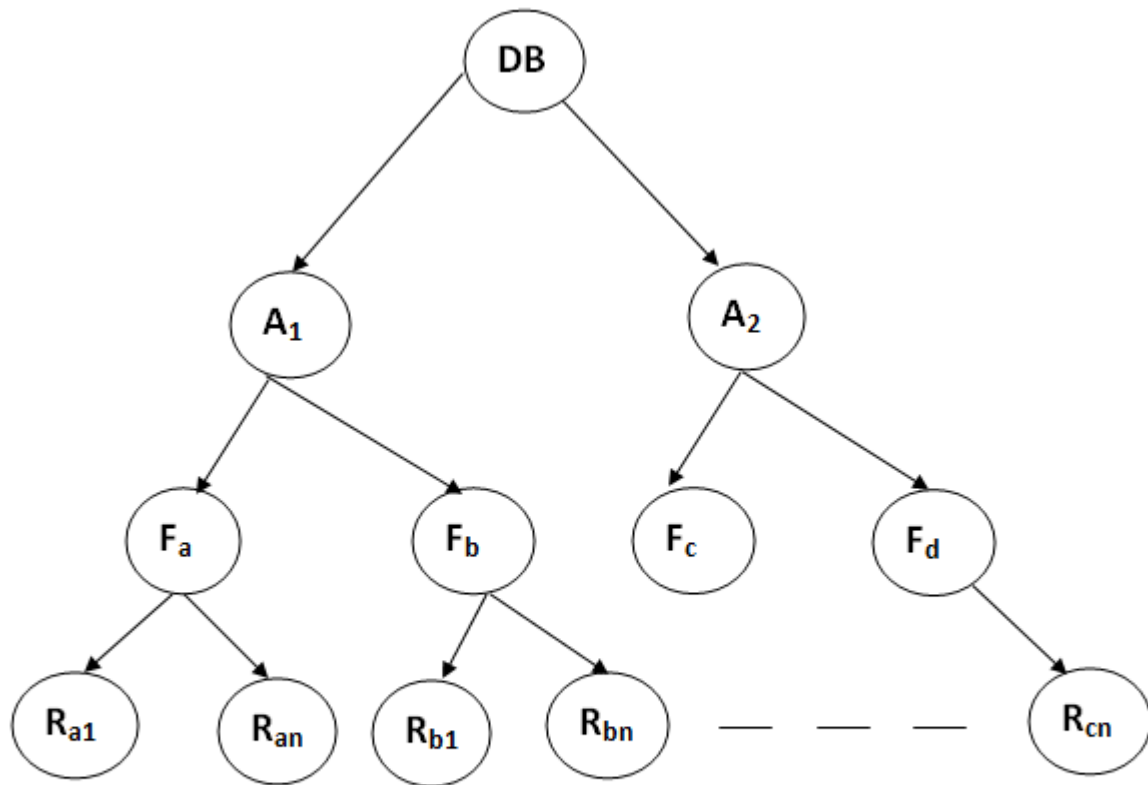
**Multiple Granularity:**

- It can be defined as hierarchically breaking up the database into blocks which can be locked.
- The Multiple Granularity protocol enhances concurrency and reduces lock overhead.
- It maintains the track of what to lock and how to lock.
- It makes easy to decide either to lock a data item or to unlock a data item. This type of hierarchy can be graphically represented as a tree.

**For example:** Consider a tree which has four levels of nodes.

- The first level or higher level shows the entire database.
- The second level represents a node of type area. The higher level database consists of exactly these areas.
- The area consists of children nodes which are known as files. No file can be present in more than one area.
- Finally, each file contains child nodes known as records. The file has exactly those records that are its child nodes. No records represent in more than one file.
- Hence, the levels of the tree starting from the top level are as follows:
  1. Database
  2. Area

3. File
4. Record



**Figure:** Multi Granularity tree Hierarchy

In this example, the highest level shows the entire database. The levels below are file, record, and fields.

There are three additional lock modes with multiple granularity:

#### Intention Mode Lock

**Intention-shared (IS):** It contains explicit locking at a lower level of the tree but only with shared locks.

**Intention-Exclusive (IX):** It contains explicit locking at a lower level with exclusive or shared locks.

**Shared & Intention-Exclusive (SIX):** In this lock, the node is locked in shared mode, and some node is locked in exclusive mode by the same transaction.

**Compatibility Matrix with Intention Lock Modes:** The below table describes the compatibility matrix for these lock modes:

	<b>IS</b>	<b>IX</b>	<b>S</b>	<b>SIX</b>	<b>X</b>
<b>IS</b>	✓	✓	✓	✓	X
<b>IX</b>	✓	✓	X	X	X
<b>S</b>	✓	X	✓	X	X
<b>SIX</b>	✓	X	X	X	X
<b>X</b>	X	X	X	X	X

It uses the intention lock modes to ensure serializability. It requires that if a transaction attempts to lock a node, then that node must follow these protocols:

- Transaction T1 should follow the lock-compatibility matrix.
- Transaction T1 firstly locks the root of the tree. It can lock it in any mode.
- If T1 currently has the parent of the node locked in either IX or IS mode, then the transaction T1 will lock a node in S or IS mode only.
- If T1 currently has the parent of the node locked in either IX or SIX modes, then the transaction T1 will lock a node in X, SIX, or IX mode only.
- If T1 has not previously unlocked any node only, then the Transaction T1 can lock a node.
- If T1 currently has none of the children of the node-locked only, then Transaction T1 will unlock a node.

Observe that in multiple-granularity, the locks are acquired in top-down order, and locks must be released in bottom-up order.

- If transaction T1 reads record  $R_{a9}$  in file  $F_a$ , then transaction T1 needs to lock the database, area  $A_1$  and file  $F_a$  in IX mode. Finally, it needs to lock  $R_{a2}$  in S mode.
- If transaction T2 modifies record  $R_{a9}$  in file  $F_a$ , then it can do so after locking the database, area  $A_1$  and file  $F_a$  in IX mode. Finally, it needs to lock the  $R_{a9}$  in X mode.
- If transaction T3 reads all the records in file  $F_a$ , then transaction T3 needs to lock the database, and area A in IS mode. At last, it needs to lock  $F_a$  in S mode.
- If transaction T4 reads the entire database, then T4 needs to lock the database in S mode.

## 10.Recovery with Concurrent Transaction

- Whenever more than one transaction is being executed, then the interleaved of logs occur. During recovery, it would become difficult for the recovery system to backtrack all logs and then start recovering.

- To ease this situation, 'checkpoint' concept is used by most DBMS.

As we have discussed **checkpoint**

in Transaction Processing Concept of this tutorial, so you can go through the concepts again to make things more clear.