

Database Creation and Management

Database Recovery System

Sasmita Kumari Nayak
CSE Department

RECOVERY MANAGEMENT

- Recovery management is a service provided by the DBMS to restore the database to a correct state after a failure.
- Here we look at the different types of failures, the tools used by recovery management and the recovery process that uses these tools.
- Database recovery is the process of restoring the database to a correct state in the event of a failure.

TYPES OF DATA BASE FAILURE

- There are many different types of failure that can affect database processing, each of which has to be dealt with in a different manner.
- Some failures affect main memory only, while others involve non-volatile (secondary) storage.
- The causes for failure are:
 - **System crashes** due to hardware, software or network errors, affecting main memory.
 - **Media failures** such as disk crash, affecting parts of secondary storage.

Cont...

- **Application software errors** such as logical errors in the program accessing the database, which cause one or more transactions to fail.
- **Natural physical disasters** such as power failures, fires, floods or earthquakes.
- **Carelessness** or unintentional destruction of data or facilities by users.
- **Sabotage**, or intentional corruption or destruction of data, hardware, or software facilities.

RECOVERY TOOLS (OR) FACILITIES

- The DBMS is equipped with the following tools to support in the recovery process.
 1. **Log File:** logging facilities, which keep track of the current state of transactions and database changes;
 2. **Checkpoint:** a checkpoint facility, which enables updates to the database that are in progress to be made permanent;
 3. **Backup mechanism:** a backup mechanism, which makes periodic backup copies of the database;
 4. **Recovery manager:** a recovery manager, which allows the system to restore the database to a consistent state following a failure.

Log File

- The file contains information about all changes made to the database.
- It provides a history of database changes made by the transactions.
- The log file contains the following information:
 - Transaction records
 - Checkpoint records

Cont...

- **Transaction records:** which contain
 - Transaction identifier
 - Type of log record (transaction start, insert, update, delete, abort, commit)
 - Identifier of data item (table, row, column) affected by the database action (insert, delete and update operations)
 - Before-image of the data item, the value before change (update and delete operations only)
 - After-image of the data item, the value after change (insert and update operations only)
 - Log management information such as pointer to previous and next log records for the transaction.

Cont...

Tid	Time	Operation	Object	Before image	After image	pPtr	nPtr
T1	10:12	START				0	2
T1	10:13	UPDATE	STAFF SL21	(old value)	(new value)	1	8
T2	10:14	START				0	4
T2	10:16	INSERT	STAFF SG37		(new value)	3	5
T2	10:17	DELETE	STAFF SA9	(old value)		4	6
T2	10:17	UPDATE	PROPERTY PG16	(old value)	(new value)	5	9
T3	10:18	START				0	11
T1	10:18	COMMIT				2	0
	10:19	CHECKPOINT	T2, T3				
T2	10:19	COMMIT				6	0
T3	10:20	INSERT	PROPERTY PG4		(new value)	7	12
T3	10:21	COMMIT				11	0

Figure 9.1: A section or segment of the log file

Cont...

- The recovery manager can perform two operations on the log. These are the:
 - **Undo** operation where the database reverts to the previous state by substituting the old value (before image) for whatever value is stored in the database (after image).
 - **Redo** operation where the database established a new state by substituting a new value (after image) for whatever value is stored in the database (before image).

Checkpoint

- The checkpoint is the time when all transactions stop.
- At this point a checkpoint record is written to the log and database buffers are written to disk.
- Checkpoints are written at periodic intervals and involve the following operations:
 - Writing all log records in main memory to secondary storage
 - Writing all modified blocks in database buffers to secondary storage
 - Writing a checkpoint record to the log file

Backup Mechanism

- At regular intervals backup copies of the database and the log file are made.
- Backups are used if the database has been damaged or destroyed.
- Backups can be complete copy of the database or incremental backup consisting only of the modifications made since the last complete or incremental backup.

Transactions and Recovery

- It is the responsibility of the recovery manager to ensure that the “atomicity” and the “durability” properties of transactions are maintained if a failure should occur.
- Let us consider the read and write operations in the database.
- For example if we want to give an employee salary rise of 5%.

Cont...

- To perform the read operation the database does the following:
 - Find the address of the disk block that contains the employee record with primary key value.
 - Transfer the disk block into a database buffer in main memory;
 - Copy the salary data from the database buffer into a variable *salary*.

Cont..

- To perform the write operation the database does the following:
 - Find the address of the disk block that contains the employee record with primary key value.
 - Transfer the disk block into a database buffer in main memory;
 - Copy the salary data from a variable *salary* into the database buffer;
 - Write the database buffer back to disk.
- Database buffers are in main memory where data is transferred to and from secondary storage.

Cont..

- Buffers are flushed to secondary storage when they are full or when the transaction commits. It is only then that update operations are considered as permanent.
- If the failure occurs between writing to the buffers and flushing the buffers to secondary storage the recovery manager must determine if the transaction has committed or not.

Cont..

- If the **transaction has committed** then to ensure durability the recovery manager has to **redo** the transaction updates to the database.
- If the **transaction has not committed** then the recovery manager has to **undo** the effects of the transaction on the database to ensure atomicity.

RECOVERY TECHNIQUES

- The type of recovery technique used is dependent on the extent of damage to the database.
- If the damage to the database is massive then the last backup copy of the database will be restored and to reapply the update operations of committed transactions using the log file.
- If the database is not physically damaged but is in a inconsistent state, then there are two techniques available to the DBMS to help the database recover. These two techniques are known as the **deferred update** and the **immediate update**.

Deferred Update

- In the deferred update approach updates are written to the database only when the database reaches the commit point.
- Updates are not written to the database at checkpoint time except for committed transactions.
- Therefore for transactions which have not been committed the undo operations are not necessary and are not used in the deferred update technique. However the recovery manager has to redo the committed transactions to ensure that the changes are permanently on the database.

Cont...

- The log records will be examined at failure to identify the active transactions at the time the failure occurred.
- Starting at the last entry in the log file, we go back to the most recent checkpoint record.
 - Any transaction with the transaction start and the transaction commit log records should be redone. The records will be redone using the after image log records for the transactions.
 - Any transaction with the transaction start and transaction abort log records nothing need be done.

Immediate Update

- In this approach updates are written to the database as they occur without waiting to reach the commit point. Database writes also occur at checkpoint time.
- If the transaction has committed then the recovery manager has to redo the transaction updates to the database.
- If the transaction has not committed then the recovery manager has to undo the effects of the transaction.
- Database writes must occur after the corresponding writes to the log file. This is known as write ahead log protocol.

Cont...

- If updates are made to the database first and the failure occurred before writing to the log records the recovery manager would not be able to identify which transaction needs to be undone or redone.
- The recovery manager examines the log file after a failure to determine if the transaction needs to be undone or redone.
 - Any transaction with the transaction start and the transaction commit log records should be redone. The records will be redone using the after image log records for the transactions.
 - Any transaction with the transaction start but not the transaction commit log records should be undone. The records will be undone using the before image log records.

Example of Deferred and Immediate Update

- To help you understand recovery from a system failure figure 9.2 shows a number of transactions with the commit time, the most recent checkpoint t_c and the failure t_f .

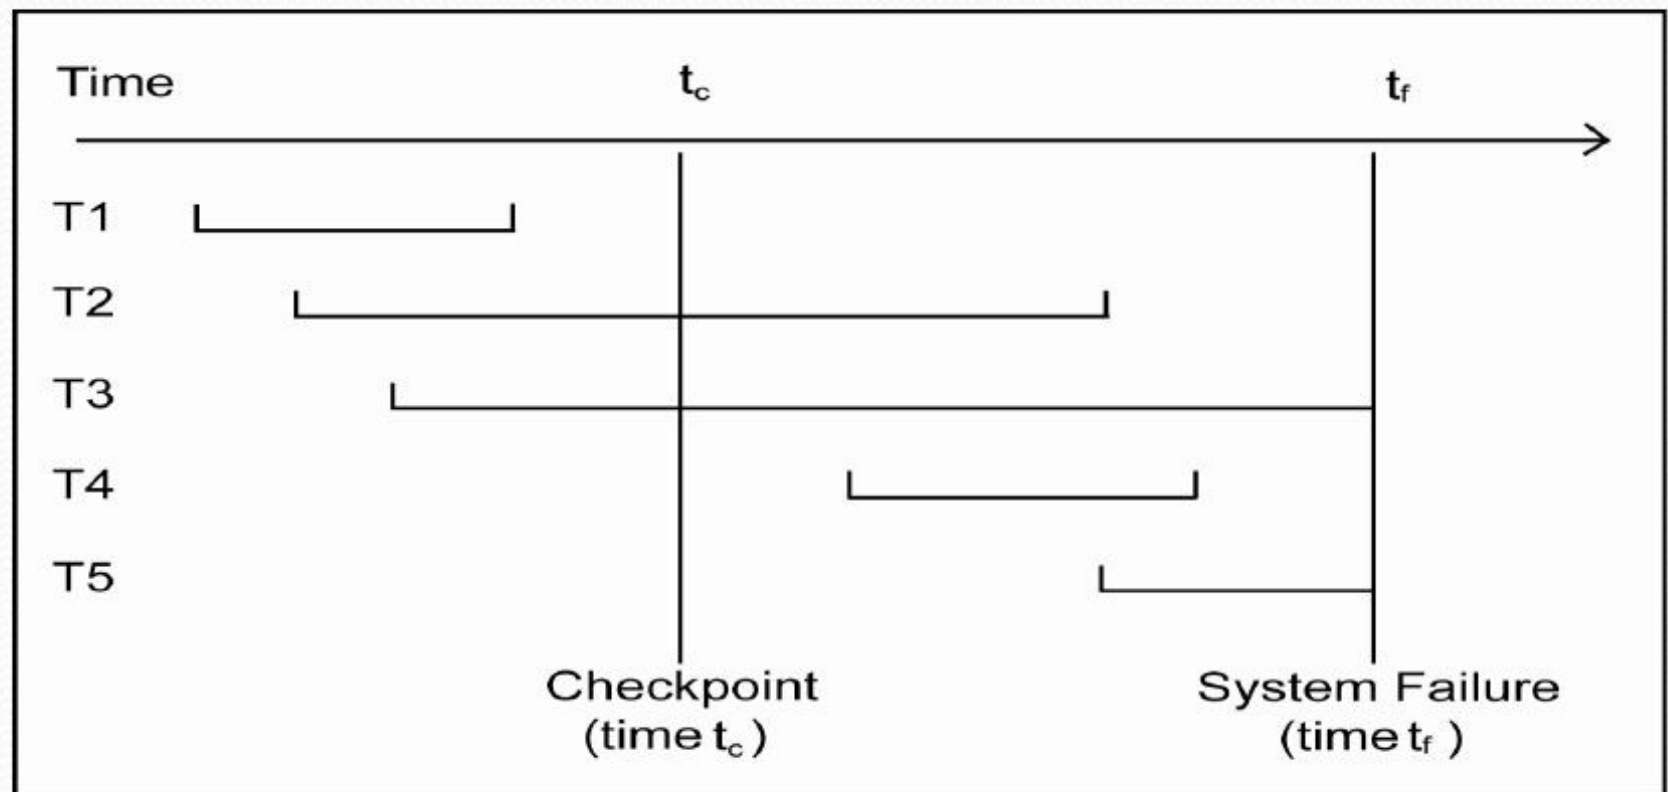


Figure 9.2: Transaction Timeline

Cont...

- A summary of recovery operations for the transaction timeline in Figure 9.2 for deferred update is shown in figure 9.3.

Transaction ID	Description	Operation
T1	Finished before CP	None
T2	Started before CP finished before failure	Redo
T3	Started before CP not yet finished	None
T4	Started after CP finished before failure	Redo
T5	Started after CP not yet finished	None

Figure 9.3: Summary of restart work for the deferred update technique

Cont...

- A summary of recovery operations for the transaction timeline in Figure 9.2 for immediate update is shown in figure 9.4.

Transaction ID	Description	Operation
T1	Finished before CP	None
T2	Started before CP finished before failure	Redo
T3	Started before CP not yet finished	Undo
T4	Started after CP finished before failure	Redo
T5	Started after CP not yet finished	Undo

Figure 9.4: Summary of restart work for the immediate update technique