



# PRESENTATION ON DATABASE RECOVERY TECHNIQUES



# INTRODUCTION

## Database Recovery

- **Pre-condition** : At any given point in time the database is in a consistent state
  - **Condition** : Some kind of **system failure** occurs
  - **Post-condition** : Restore the database to the consistent state that existed before the failure
- 
- **Database recovery** is the process of restoring the database to the most recent consistent state that existed just before the failure
  - **Example**:
    - If the system crashes before a fund transfer transaction completes its execution, then either one or both accounts may have incorrect value. Thus, the database must be restored to the state before the transaction modified any of the accounts
- 

# FAILURE CLASSIFICATION

- Types of failures
  - The database may become unavailable for use due to
    - **Transaction failure**: Transactions may fail because of incorrect input, deadlock, incorrect synchronization
    - **System failure**: System may fail because of addressing error, application error, operating system fault, RAM failure, etc
    - **Media failure**: Disk head crash, power disruption, etc



# TRANSACTION LOG

- For recovery from any type of failure data values prior to modification (BFIM - BeFore Image) and the new value after modification (AFIM – After Image) are required
- These values and other information is stored in a sequential file called Transaction log. A sample log is given below. Back P and Next P point to the previous and next log records of the same transaction

# MAIN RECOVERY TECHNIQUES

## 1. Deferred Update:

- **These techniques do not physically update the DB on disk until after a transaction reaches its commit point**
- **The updates are recorded in the local transaction buffer and in the log file for recovery**
- **These techniques need only to redo the committed transaction and no-undo is needed in case of failure (No-Undo/Redo)**



# DEFERRED UPDATE EXAMPLE IN A SINGLE USER ENVIRONMENT

(a)

$T_1$	$T_2$
read_item( $A$ )	read_item( $B$ )
read_item( $D$ )	write_item( $B$ )
write_item( $D$ )	read_item( $D$ )
	write_item( $D$ )

(b)

[start_transaction, $T_1$ ]
[write_item, $T_1, D, 20$ ]
[commit, $T_1$ ]
[start_transaction, $T_2$ ]
[write_item, $T_2, B, 10$ ]
[write_item, $T_2, D, 25$ ]

← System crash

**Figure 19.2**

An example of recovery using deferred update in a single-user environment. (a) The READ and WRITE operations of two transactions. (b) The system log at the point of crash.

The [write\_item,...] operations of  $T_1$  are redone.

$T_2$  log entries are ignored by the recovery process.



# MAIN RECOVERY TECHNIQUES(CONT)

## 2. Immediate Update:

- **The DB may be updated by some operations of a transaction before the transaction reaches its commit point**
- **The updates are recorded in the log which must contain the old values (BFIM) and the new values (AFIM)**
- **These techniques need to undo the operations of the uncommitted transactions and redo the operations of the committed transactions (Undo/Redo)**
- **The Undo/No-Redo may be used in special case where all operations are recorded in the DB before the transaction commits**



# MAIN RECOVERY TECHNIQUES(CONT)

## 3. Shadow Update:

- The modified version of a data item does not overwrite its disk copy but is written at a separate disk location

## 4. In-Place Update:

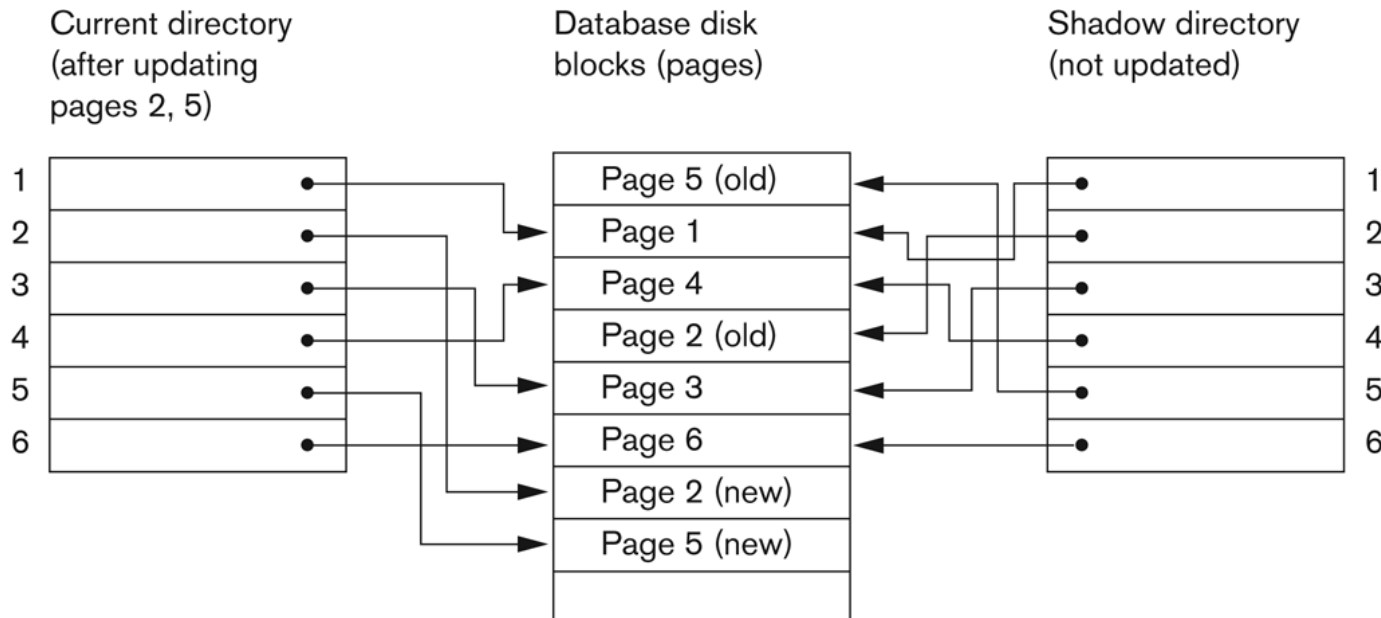
- The disk version of the data item is overwritten by the cache version





# MAIN RECOVERY TECHNIQUES(CONT)

## Example of Shadow paging:



**Figure 19.5**  
An example of shadow paging.



# CHECKPOINTS

## Checkpointing:

- Time to time (randomly or under some criteria) the database flushes its buffer to database disk to minimize the task of recovery. The following steps defines a checkpoint operation:
  1. Suspend execution of transactions temporarily
  2. Force write modified buffer data to disk
  3. Write a [checkpoint] record to the log, save the log to disk.
  4. Resume normal transaction execution
- During recovery redo or undo is required to transactions appearing after [checkpoint] record

# RECOVERY TECHNIQUES BASED ON DEFERRED UPDATE(No UNDO/REDO)

- The idea is to postpone any actual updates to the DB until the transaction completes its execution and commits (it follows the No-Steal approach)
- The updates are recorded only in the log and in the buffer
- After transaction reaches its commit point and the log is force-written to disk, the updates are recorded in the disk
- If the transaction fails before commit, no need to undo any operations



# RECOVERY TECHNIQUES BASED ON IMMEDIATE UPDATE(UNDO/NO-REDO)

- In this algorithm AFIMs of a transaction are flushed to the database disk under WAL before it commits
  - For this reason the recovery manager **undoes** all transactions during recovery
  - No transaction is **redone**
  - It is possible that a transaction might have completed execution and ready to commit but this transaction is also **undone**

