

CS 5300 module7

Name

Student ID

Problem #1 (15 Points)

Shadow paging is an alternative to log-based recovery procedure:

- a) What is it and how does it work (I need clear and detail explanation)?

Database is partitioned into fixed-length blocks (i.e., pages). A page table is used to hold addresses of pages on the disk. Each page has an entry in the page table.

In shadow paging, two page tables are going to be maintained during the course of a transaction: Current Page Table and Shadow page table

At the start of a transaction, both page tables are identical. During the course of transaction, current page table may change (due to write operation), but shadow page table remains unchanged.

All input and output operations use the current page table to locate database pages on disk.

When a transaction commits, the current page table becomes the new shadow page table and the next transaction is allowed to start execution. Critical issue is the fact that the shadow page table must be stored in a non-volatile storage since it provides the only means of locating database pages.

After system comes back up, it copies the shadow page table into main memory and uses it for subsequent transactions. Unlike log-based schemes, it does not need to invoke undo operations

- b) Discuss about advantages and disadvantages of shadow paging (relative to the log-based protocols).

Under certain circumstances, shadow paging offers fewer disk accesses. However, it is much harder to be extended for concurrent execution of transactions.

No Undo/No Redo algorithm

Data locality will be lost (fragmentation)

Garbage collection overhead

Commit overhead

Too many page copies (problem if dealing with large databases and lots of changes).