

Question 1: Buffer Management.....[390 points]

- (i) [10 points] **Buffer Management:**
Distinguish between disk management and buffer management.
- (ii) [10 points] **Locality of Reference:**
Define locality of reference: (1) spatial and (2) temporal.
- (iii) [10 points] **Buffer Manager:**
Define the input and output of a page table.
- (iv) [10 points] **Buffer Manager:**
List the purpose of these meta-data: (1) dirty flage, and (2) reference counter.
- (v) [30 points] **Locks vs Latches:**
List three differences between locks and latches.
- (vi) [10 points] **Locks vs Latches:**
With latches, we do not need to be able to rollback changes. Why?
- (vii) [10 points] **Page Table vs. Page Directory:**
Distinguish between page table and page directory.
- (viii) [10 points] **Page Table vs. Page Directory:**
Distinguish between page table and page directory.
- (ix) [10 points] **Buffer Manager Interface:**
List the two key functions supported by the buffer manager.
- (x) [10 points] **Buffer Manager:**
What data structure is used to implement a page table inside the buffer manager?
- (xi) [10 points] **Buffer Manager:**
What is the purpose of page-level latch?
- (xii) [10 points] **Logging:**
What is the log sequence number (LSN) of a page?
- (xiii) [10 points] **Logging:**
What is write-ahead logging (WAL)? Justify its name.
- (xiv) [30 points] **Logging:**
To update a tuple in a table, list the three operations done in a DBMS that uses write-ahead logging. Why are these operations needed?
- (xv) [10 points] **Multiple Buffer Pools:**
List a benefit and a limitation of maintaining multiple buffer pools.
- (xvi) [10 points] **Pre-fetching:**
Explain pre-fetching in the context of sequential heap scan with an example.
- (xvii) [10 points] **Pre-fetching:**
Explain pre-fetching in the context of index scan with an example.
- (xviii) [10 points] **Scan Sharing:**
Define scan sharing. Does it require spatial locality or temporal locality or both?

- (xix) **[10 points] Buffer Pool Bypass:**
Define buffer pool bypass optimization. How is it related to scan-resistant buffer replacement policies?
- (xx) **[10 points] Direct I/O:**
Define direct I/O.
- (xxi) **[10 points] OS Page Cache:**
List two limitations of relying on the OS for moving data between memory and disk.
- (xxii) **[10 points] Background Writing:**
Define background writing optimization. How does the DBMS ensure the atomicity property?
- (xxiii) **[10 points] ACID:**
Define the ACID properties. Which two properties are guaranteed by the WAL protocol?
- (xxiv) **[10 points] Other Memory Pools:**
What is anonymous mapping? Distinguish it from file-backed mapping. What are the flags passed to `mmap` for these two types of memory mapping?
- (xxv) **[10 points] Buffer Replacement Policies:**
Define the buffer replacement problem. What is the optimal policy?
- (xxvi) **[10 points] Buffer Replacement Policies:**
What is a page fault? How does it vary with buffer size?
- (xxvii) **[10 points] Policy vs Mechanism:**
Distinguish between policy and mechanism with an example.
- (xxviii) **[10 points] FIFO:**
List a limitation of the FIFO policy.
- (xxix) **[10 points] LRU vs LFU:**
When does LRU outperform LFU? When does LFU outperform LRU?
- (xxx) **[10 points] CLOCK:**
Why does the CLOCK policy approximate LRU?
- (xxxi) **[10 points] Sequential Flooding:**
Define the sequential flooding problem.
- (xxxii) **[10 points] Sequential Flooding:**
Why is LRU-K resistant to the sequential flooding problem?
- (xxxiii) **[10 points] Sequential Flooding:**
Why is 2Q resistant to the sequential flooding problem?
- (xxxiv) **[10 points] 2Q vs LRU-K:**
How is 2Q related to LRU-K?
- (xxxv) **[10 points] Priority Hints:**
Define priority hints. Illustrate with an example.