Question 1: Concurrency Control in MMDBs . . . . . [150 points]

(i) **[10 points] CC in Main-Memory DBMSs:**
Distinguish between CC in: (1) disk-oriend DBMSs and (2) main-memory DBMSs.

(ii) **[10 points] CC in Main-Memory DBMSs:**
Explain how you would use the compare-and-swap instruction in a CC protocol.

(iii) **[10 points] CC in Main-Memory DBMSs:**
How does 2PL function under high contention?

(iv) **[10 points] CC in Main-Memory DBMSs:**
How does TO function under high contention?

(v) **[10 points] CPU Simulator:**
Distinguish between: (1) NUMA and (2) NUCA architectures.

(vi) **[10 points] Query Types:**
Distinguish between: (1) stored procedure, (2) prepared statement, and (2) ad-hoc query.

(vii) **[10 points] Workload Types:**
Distinguish between: (1) YCSB, (2) TPC-C, and (3) TPC-H workloads.

(viii) **[10 points] CC Protocol Comparison:**
Which CC protocol scales well on a read-only workload? Why?

(ix) **[10 points] CC Protocol Comparison:**
Which CC protocol scales well on a write-intensive/medium-contention workload? Why?

(x) **[10 points] CC Protocol Comparison:**
Which CC protocol scales well on a write-intensive/high-contention workload? Why?

(xi) **[10 points] CC Protocol Comparison:**
What are the bottlenecks in CC protocols under high-contention?

(xii) **[10 points] Lock Thrashing:**
Define lock thrashing. Suggest a way to reduce the impact of lock thrashing.

(xiii) **[10 points] Timestamp Allocation:**
Explain how batched atomic addition helps reduce contention.

(xiv) **[10 points] Memory Allocation:**
Explain why the default libc malloc is slow under high contention.

(xv) **[10 points] Hardware Trend:**
Explain why increases in DRAM capacities have stalled in recent years compared to SSDs.

End of Problem Set