Question 1: Persistent Memory Databases ..... [330 points]

(i) **[10 points] Disk-oriented vs In-Memory DBMSs:**
List two differences between in-memory DBMS and a disk-oriented DBMS with a large buffer pool.

(ii) **[10 points] Device Latency:**
What is the approximate latency of: (1) SRAM, (2) DRAM, (3) NVM, (4) SSD, and (5) HDD.

(iii) **[10 points] Memory Pools:**
Justify the need for separate memory pools for fixed-length and variable-length data.

(iv) **[10 points] Importance of Hardware:**
Define Moore's Law. How did it influence the adoption of database machines?

(v) **[10 points] Importance of Hardware:**
Define: (1) FPGA, (2) GPU, and (3) Configurable Spatial Accelerator (CSA).

(vi) **[10 points] Persistent Memory:**
Distinguish between: (1) NVM and DRAM, (2) NVM and SSD.

(vii) **[10 points] Interfaces:**
How is NVMe related to PCIe?

(viii) **[10 points] Persistent Memory:**
Define Direct Access (DAX) feature of NVM-aware file systems.

(ix) **[10 points] Persistent Memory:**
Distinguish between: (1) CLFLUSH and (2) CLWB instructions.

(x) **[10 points] Persistent Memory:**
Distinguish between: (1) Memory Mode and (2) App Direct Mode. Which mode is better suited for an NVM-aware DBMS?

(xi) **[10 points] Persistent Memory:**
Why are in-memory DBMSs better positioned to use byte-addressable NVM as opposed to disk-oriented DBMSs?

(xii) **[10 points] Synchronization:**
List the two assembly instructions used in the synchronization primitive.

(xiii) **[10 points] Synchronization:**
Explain the ADR feature.

(xiv) **[10 points] Naming:**
Distinguish between a volatile pointer and a persistent pointer.

(xv) **[10 points] PM-Aware Allocator:**
Distinguish between a regular allocator and a PM-aware allocator.

(xvi) **[10 points] Storage Engine Architectures:**
List three canonical storage engine architectures.

Question 1 continues...

(xvii) **[10 points] Storage Engine Architectures:**
List the write operations in an in-place updates engine for an INSERT statement.

(xviii) **[10 points] Storage Engine Architectures:**
List the benefits and limitations of the in-place updates engine.

(xix) **[10 points] Storage Engine Architectures:**
List the write operations in an NVM-aware in-place updates engine for an INSERT statement.

(xx) **[10 points] Storage Engine Architectures:**
List the benefits and limitations of the copy-on-write engine.

(xxi) **[10 points] Storage Engine Architectures:**
Distinguish between an in-place updates engine and a copy-on-write engine.

(xxii) **[10 points] Storage Engine Architectures:**
Justify the name – "copy-on-write" or "shadow paging" engine.

(xxiii) **[20 points] Storage Engine Architectures:**
Is an NVM-aware copy-on-write engine different from an NVM-aware in-place updates engine from a recovery latency standpoint?

(xxiv) **[20 points] Storage Engine Architectures:**
List the write operations in an log-structured updates engine for an INSERT statement.

(xxv) **[10 points] Storage Engine Architectures:**
List the benefits and limitations of the log-structured updates engine.

(xxvi) **[20 points] Storage Engine Architectures:**
List the write operations in an NVM-aware log-structured updates engine for an INSERT statement.

(xxvii) **[10 points] Storage Engine Architectures:**
How is a tuple stitched together in a log-structured updates engine?

(xxviii) **[10 points] Storage Engine Architectures:**
What is the purpose of compaction in a log-structured updates engine?

(xxix) **[10 points] Storage Engine Architectures:**
What is the purpose of a Bloom Filter in a log-structured updates engine?

(xxx) **[10 points] Storage Engine Architectures:**
What is the purpose of optimizing storage engine architectures for NVM?