

Question 1: Data Representation [190 points]

- (i) [10 points] **Data Representation:**
Define a tuple schema. Why is it needed in a relational DBMS?
- (ii) [10 points] **Numeric Data:**
Distinguish between fixed and variable precision numbers.
- (iii) [10 points] **Numeric Data:**
Distinguish between FLOAT and NUMERIC data types. Which one is faster to evaluate? Why?
- (iv) [10 points] **Numeric Data:**
If your application cannot handle rounding errors, which data type should you use for storing numeric data?
- (v) [10 points] **Numeric Data:**
What is the purpose of scale in NUMERIC data type?
- (vi) [10 points] **Large Values:**
How does a DBMS store values larger than a page?
- (vii) [10 points] **External Value Storage:**
List two limitations of external value storage.
- (viii) [20 points] **External Value Storage:**
Why are small objects (< 256 KB) best stored in the DBMS? Why are large objects (> 1 MB) best stored in the filesystem?
- (ix) [15 points] **Workloads:**
Distinguish between: (1) OLTP, (2) OLAP, and (3) HTAP workloads with respect to complexity of operations (*i.e.* number of tuples affected by the operation).
- (x) [15 points] **Workloads:**
Distinguish between: (1) OLTP, (2) OLAP, and (3) HTAP workloads with respect to type of operations (*i.e.* read or write).
- (xi) [15 points] **Storage Models:**
Distinguish between: (1) OLTP, (2) OLAP, and (3) HTAP workloads with respect to the optimal storage model.
- (xii) [15 points] **Storage Models:**
Distinguish between: (1) NSM and (2) DSM.
- (xiii) [15 points] **Storage Models:**
Why does NSM not work well for OLAP workloads?
- (xiv) [15 points] **Decomposition Storage Model:**
When are embedded tuple ids needed in case of DSM?
- (xv) [10 points] **Schema Changes:**
Is ADD COLUMN faster to execute with NSM or DSM? Justify your answer.