

Question 1: Scheduling [240 points]

- (i) **[10 points] Query Execution:**
Define these terms: (1) operator, (2) task, and (3) pipeline.
- (ii) **[10 points] Data Placement:**
Distinguish between uniform and non-uniform memory access (NUMA) layouts.
- (iii) **[10 points] Data Placement:**
Explain the connection between distributed DBMSs and a single-node NUMA layout.
- (iv) **[10 points] Data Placement:**
List two ways to use `numactl`.
- (v) **[10 points] Memory Allocation:**
Justify the term "virtual memory".
- (vi) **[10 points] Memory Allocation:**
List the system call used by the OS to implement `malloc`.
- (vii) **[10 points] Memory Allocation:**
Distinguish between these two policies: interleaving and first-touch.
- (viii) **[10 points] Data Placement:**
Explain the significance of data placement in: (1) OLTP and (2) OLAP workloads.
- (ix) **[10 points] Data Placement:**
Distinguish between partitioning and placement schemes.
- (x) **[10 points] Worker Allocation:**
Define these terms: (1) worker, and (2) core.
- (xi) **[10 points] Worker Allocation:**
List two policies for allocating workers.
- (xii) **[10 points] Worker Allocation:**
List two techniques for assigning tasks.
- (xiii) **[10 points] Worker Allocation:**
List a benefit and a drawback of using a centralized dispatcher.
- (xiv) **[10 points] Scheduling:**
Distinguish between static and dynamic scheduling.
- (xv) **[10 points] Scheduling:**
Define a straggler. List a technique to cope with stragglers.
- (xvi) **[10 points] Scheduling:**
Define a morsel.
- (xvii) **[10 points] Scheduling:**
How can we reduce cross-communication between workers?
- (xviii) **[10 points] Scheduling:**
Define a watchdog thread.

- (xix) **[10 points] Scheduling:**
Distinguish between soft and hard queues.
- (xx) **[10 points] Scheduling:**
Why is work-stealing not beneficial for systems with a larger number of sockets?
- (xxi) **[10 points] Scheduling:**
Distinguish between preemptive and non-preemptive thread scheduling.
- (xxii) **[10 points] Scheduling:**
Explain how SQLOS operates in user-mode.
- (xxiii) **[10 points] Scheduling:**
Define a quantum. Why is it hard to enforce this scheduling constraint?
- (xxiv) **[10 points] Flow Control:**
List two techniques for flow control.