
Exercise 1 - Throughput and Response Time.**2 Point**

A database system must process transactions t_1 , t_2 , t_3 , and t_4 . For each transaction, Table 1 shows the start time (when the transaction enters the system) and the execution time (the runtime of the transaction in the system when it is not interrupted). All time values are in seconds.

Compute the average response time and throughput for the following task scheduling strategies:

1. The transactions are executed in the order of arrival, i.e., the waiting queue is ordered by start time.
2. Shorter transactions are given priority over longer transactions, i.e., the waiting queue is ordered by execution time. A running transaction is aborted and put back to the waiting queue for re-execution if a shorter transaction (with a smaller execution time) enters the system.

Transaction	Start Time	Execution Time
t_1	0	4
t_2	1	1
t_3	2	2
t_4	3	3

Tabelle 1: Start and execution time of transactions.

*Exercise 2 - Parallel System Architecture.*2 Point

Mark the following statements as true (**T**) or false (**F**).

1. The shared memory and the shared disk architecture are identical.
2. NUMA allows processors to access the memory of other processors through a high-speed network.
3. Shared memory scales to larger number of processors than shared nothing architectures.
4. The top level of hierarchical architectures is shared nothing.

Exercise 3 - Horizontal Partitioning.2 Point

Relation $r[A]$ in Table 2 should be horizontally partitioned onto four disks, D_i , $0 \leq i \leq 3$. Partition the tuples on attribute A using

1. round-robin (with processing order left-to-right in Table 2), and
2. hash partitioning with hash function $h(a) := a \bmod 4$.

Further, answer the following questions:

3. What is the downside of hash partitioning compared to round-robin?
4. For which type of queries is hash partitioning preferable over round-robin?

A	30	72	54	46	66	34	42	60	10	22	84	96
-----	----	----	----	----	----	----	----	----	----	----	----	----

Tabelle 2: Relation $r[A]$.

Exercise 4 - 2-Phase-Commit (2PC).**2 Point**

Transaction T is initiated at site S_i with coordinator C_i , $1 \leq i \leq n$, and is executed at sites S_k , $1 \leq k \leq n$. Discuss the situations when a message between coordinator C_i and site S_k gets lost in

1. Phase 1 of the protocol, and
2. Phase 2 of the protocol.

Exercise 5**1 Point**

Discuss the persistent messaging protocol between a sender S and a receiver R for the following situations:

1. Message M_1 is received by R , R sends an acknowledgement to S , which S receives.
2. Message M_2 is received by R , R sends an acknowledgement to S , but S does not receive the acknowledgement.
3. Message M_3 is sent multiple times by S .
4. Message M_4 was sent multiple times. One of the messages was delayed in the network and arrives after R has deleted M_4 from its received messages relation.

Exercise 6

1 Point

Consider the following transactions on the replicated data items X and Y , initially $X = 5$ and $Y = 10$ on all replicas.

$$T_1 : R(Y), R(X) \quad T_2 : W(X \leftarrow 8) \quad T_3 : R(X), Y \leftarrow X, W(Y)$$

1. What are the possible values T_1 reads for X and Y under strong consistency?
2. Show how T_1 could read values for X and Y that cannot result from any strongly consistent execution of the transactions if the replication protocol does not enforce strong consistency.

Exercise 7 - Deadlock Handling.**2 Point**

Consider three transactions $T_1 : W(X), W(Y), W(Z)$, $T_2 : W(Y), W(Z), W(X)$, and $T_3 : W(Z), W(X), W(Y)$ that update data items X , Y , and Z .

1. Provide a schedule using 2-phase locking that results in a cycle that includes all transactions.
2. Draw the according local and the global wait-for graphs assuming a distributed system and a distributed lock manager. Data item X is managed by site S_1 , Y is managed by site S_2 , and Z is managed by site S_3 .
3. Explain how the cycle is resolved in the centralized approach.

Exercise 8 - *Range-Partitioning Sort.***2 Point**

Execute range-partitioning sort on 20 values in the range 0 ... 99 that are horizontally partitioned (round robin) on five processors, P_i , $0 \leq i \leq 4$, in a shared-nothing environment. The processors P'_0 , P'_1 , and P'_2 support the sorting.