Department of Computer Science

Prof. Dr. Nikolaus Augsten
Jakob-Haringer-Str. 2
5020 Salzburg, Austria
Telefon: +43 662 8044 6347
E-Mail: nikolaus.augsten@plus.ac.at

Paris Lodron Universität Salzburg

Non-Standard Database Systems — Exam
Summer Semester 2022/2023 — 19.09.2023

Name: ____________________ **Student ID:** ____________________

Hints

- Check whether you received all pages of the exam (9 pages).
- Write your name or your student ID on each sheet of the exam and hand in all pages.
- All answers are expected to be written on the exam sheet.
- Number any additional pages that you use for longer answers and reference them clearly visible on the corresponding exercise sheet.
- Use only pencils that are permanent. No red pencil.
- Use the notation and techniques discussed during the lecture.
- Exercises with more than one solution are not graded.
- You may use one A4 sheet with your personal notes (on both sides, hand written or printed).
- Exam duration: 90 minutes

Signature ____________________

Grading

Filled by the examiner

Exercise	1	2	3	4	5	6	7	8	Summe
Total Points	2	2	2	2	2	2	2	2	16
Points Reached									

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

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).

Compute the average response time and the throughput for the following task scheduling strategies where shorter transactions are given priority over longer transactions (i.e., the waiting queue is ordered by execution time):

1. No running transaction is aborted.
2. Abort the running transaction when a shorter transaction enters the waiting queue.

Transaction	Start Time	Execution Time
t_1	0	7
t_2	2	1
t_3	4	2
t_4	6	1

Table 1: Start and execution time of transactions.

Exercise 2 - *Parallel System Architecture*. 2 Points

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

1. In a parallel database system, compute nodes are located at different sites. ☐
2. The latency in WANs is higher than in LANs. ☐
3. In a mesh network architecture, all components are directly connected to each other. ☐
4. Deadlocks are detected by the lock manager. ☐

Exercise 3 - *Horizontal Partitioning.* 2 Points

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. range partitioning with partitioning vector $v = [25, 50, 75]$.

Further, answer the following question:

3. What is the downside of range partitioning?

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

Table 2: Relation $r[A]$.

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

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 following situations:

1. Site S_k ($k \neq i$) violates the 2PC protocol and sends the `ready T` message before writing all log records to stable storage.
2. Site S_k ($k \neq i$) violates the 2PC protocol and releases the locks before the end of Phase 2.

Exercise 5 - *Persistent Messaging.* 2 Points

Consider a sender S that sends a message to receiver R using the persistent messaging protocol. Table 3 shows the initial entries in the relations *messages_to_send* of the sender and *received_messages* of the receiver. Newer events have larger time stamps.

1. Compute the value of T_{OLD} for the relations in Table 3.
2. Show the relation *received_messages* after receiver R has received and processes the value of T_{OLD}.

messages_to_send

number	message	time	ack
1	$Q \leftarrow Q + 9$	5	received
3	$A \leftarrow A + 3$	9	received
7	$Q \leftarrow Q + 3$	13	
8	$B \leftarrow B - 9$	15	received
9	$C \leftarrow C - 6$	22	

received_messages

number	message	time	ack
1	$Q \leftarrow Q + 9$	5	sent
3	$A \leftarrow A + 3$	9	sent
7	$Q \leftarrow Q + 3$	13	sent
8	$B \leftarrow B - 9$	15	sent

Table 3: Relations *messages_to_send* at sender S and *received_messages* at receiver R.

Exercise 6 - *Deadlock Handling*. 2 Points

Draw the global wait-for graph for the following local wait-for graphs for sites S_1, S_2, and S_3. Further, highlight all cycles in the global wait-for graph, if applicable.

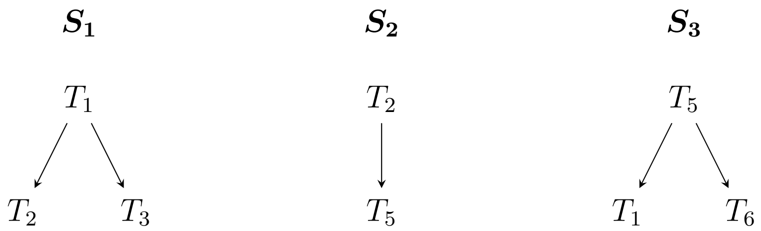

Exercise 7 - *Vector clocks.* 2 Points

Assume a single data item Q that is replicated on sites S_1, S_2, and S_3. A site S_i can do (i) a local write on Q, $W(Q)$, which changes the value of the local copy Q_i, or (ii) copy the value from a different site S_j, $C(S_j)$, $j \neq i$, which copies the value of Q_j to Q_i. All vectors are initialized with the zero vector.

1. Show the vector clocks resulting from the schedule in Figure 1.
2. Identify the first operation that leads to a branching history.

Note: Local reads, which will typically precede a local write in a real schedule, are not relevant for conflict detection and omitted from the schedule.

S_1	S_2	S_3
	$C(Q_1)$	
$W(Q)$		
		$W(Q)$
$W(Q)$		
	$C(Q_3)$	
	$C(Q_1)$	
		$C(Q_2)$

Figure 1: Schedule on replicated data item Q.

Exercise 8 - *Parallel Join.* 2 Points

Given a system with 6 processing nodes p_i, $1 \leq i \leq 6$, and three relations $R[A] = [4, 7, 13, 14, 16, 24, 25, 36, 44, 55, 57, 62, 68, 72, 78, 81, 85, 92]$, $S[A] = [7, 34]$, and $T[A] = [24, 62]$ (each number is an attribute value of A and forms a unary tuple). The tuples of the relations are round-robin partitioned on the processing nodes p_i.

Find an appropriate parallel join technique for the following query and count:

1. the number of tuples that must be transferred over the network between any pair of nodes;
2. the number of tuple pairs from R, S, and T that must be joined per processing node.

```
SELECT * FROM R, S, T
WHERE R.A <= S.A AND R.A <= T.A;
```