

The minimum for active attendance: 3 tasks done

1. Explain two essentially different ways to execute the following query. (What kind of indexes it is possible to use?)

```
select * from EMPLOYEE where SALARY < 30000
or BDATE > '20.04.1932 and BDATE < '20.04.1952';
```

2. Let us assume the following query:

```
select LNAME, FNAME from EMPLOYEE where exists
(select DNAME from DEPARTMENT where MGRSSN = EMPLOYEE.SSN);
```

- What is the meaning of this query?
- Explain how the query can be efficiently executed.

3. The following query contains a subquery:

```
select * from account where custno in
(select custno from customers)
```

Under what conditions it is possible to change it to the following query which is based on a join?

```
select accounts.* from accounts, customers
where accounts.custno = customers.custno;
```

4. a) What kind of an operation sequence is included in a typical online banking terminal transaction (a cash withdrawal)? We assume that the card's verification code is checked locally (without a database) but the accounts are managed by a database. Which exceptions there can happen (concerning the transaction management)?

b) Suppose that a system makes reservations for several successive flights (from A to B with some transfers between) or for a holiday trip (with hotel, flights, and car rental). Are these situations somehow problematic from the transaction management point of view? (Consider the ACID properties, EN, p. 640, 1-4.)

5. Assume that the following records are found in the log file when the system crashes:

```
99: [checkpoint]
100: [start, T1]
101: [start, T2]
102: [write, T1, A, 10, 20]
103: [write, T2, B, 11, 22]
104: [commit, T1]
105: [start, T3]
106: [write, T3, A, 20, 30]
107: [abort, T2]
108: [checkpoint]
109: [start, T4]
110: [write, T3, B, 11, 33]
111: [write, T3, A, 30, 40]
112: [write, T4, C, 15, 25]
113: [commit, T3]
114: [write, T4, C, 25, 35]
```

a) What recovery operations are needed? We assume that the PageLSN (page log sequence number) field of the page containing data item A is 106, the PageLSN for B is 110 and the PageLSN for C is 112.

b) The same question for the log without line 113.

6. What kind of a commit protocol should be used were it required that no redos be necessary in recovery? What drawbacks there are with this scheme?