

Group A

Please fill in your name and registration number (Matrikelnr.) **immediately**.

Exam 1 ON		SOLUTION KEY	02.05.2019
		Advanced Database Systems (184.780)	GROUP A
Matrikelnr.	Last Name	First Name	

Duration: 60 minutes. Provide the solutions on the designated pages. **Good Luck!**

Question 1: (4)

For each of the statements below, decide if it is true or false and tick the corresponding circle. You get +1 credit for each correct answer, −1 credit for each wrong answer and 0 credit if you leave the answer open. In total, you always get ≥ 0 credits on the entire exam question 1.

1. Suppose that there exists a hash index with static hashing on some integer attribute A . Then it always suffices to read 1 page in order to check if a tuple satisfying the condition $A = 17$ exists.
2. If there are m buffer frames available for sorting, then the DBMS can combine up to $m + 1$ runs in each pass of the merge phase.
3. For all relations R and S and for all selection predicates P , the following equalities hold:
 $\sigma_P(R \cup S) = \sigma_P(R) \cup \sigma_P(S)$ and $\sigma_P(R \cap S) = \sigma_P(R) \cap \sigma_P(S)$.
4. For all relations R and S and for all selection predicates P , the following equality holds:
 $\sigma_P(R \bowtie S) = \sigma_P(R) \bowtie \sigma_P(S)$.

Question 2:

Suppose that the database of a car rental company contains the following relations:

Cars(Id, Brand, Category, KW, Price) (short *c*)

Assistants(Id, Name, Salary, Address) (short *a*)

Rentings(Id, CarId, CustomerId, AssistantId, BeginDate, EndDate) (short *r*).

Further, suppose that the following query has to be processed:

```

SELECT a.Name, a.Salary, c.Brand, c.Price
FROM Cars c, Assistants a, Rentings r
WHERE c.Price > 150
AND a.Salary < 1200
AND c.Id = r.CarId
AND a.Id = r.AssistantId
  
```

i.e., we are interested in information on assistants who let expensive cars and whose salary is rather low.

a) In the first box below, draw the query tree of the canonical translation. To save space, you may use the abbreviations *c*, *a*, and *r*, for the relations Cars, Assistants, and Rentings, respectively.

b) In the second box below, draw the query tree of the optimized Relational Algebra expression. For the optimization, apply the following rules:

- Push selections as deep in the tree as possible,
- project out attributes as soon as they are not needed anymore,
- replace cross products by joins.

(a) Canonical Translation:

(b) Optimized Expression:

Question 3:

(3)

Consider a magnetic disk with the following characteristics:

- Average seek time: 3ms
- Track-to-track seek time: 0.2ms
- Transfer time (for 1 block): 1ms
- Rotational speed: 5000 rpm
- Block size: 15 000 byte
- Average track size 450 000 byte

Assume a file **books** with 1400 fixed-length records with a record size of 2000 bytes is stored on the disk in an unspanned organization.

a) Compute the the time required to read the whole **books** file from disk if you can make no assumptions on the layout of its blocks on the disk.

b) Now, assume the file **books** is stored on consecutive blocks of the disk. Compute the time required to read the whole **books** file from disk again under these new assumptions. (For simplicity you can also assume that no rotational delay occurs during track changes.)

c) There always exists a percentage p such that when reading $> p\%$ of a file with random access it is faster to read the entire file instead. Estimate this percentage for the given scenario. (You do not need to calculate it exactly, a rough calculation is sufficient.)

Question 4:

(5)

a) Your DBMS maintains equi-depth histograms to support selectivity calculation for query planning. In particular, for the column **credits** of **integers** of a table **students** with 5000 rows, the following 4 values divide the column values into 5 groups of equal size: 34, 61, 181, 320. Furthermore, the maximum value present in the column **credits** is known to be 500. **Assume that the dividing values in the histogram are always part of the left bucket, i.e., the value 34 is in the first bucket, 61 in the second, and so on.**

Estimate the selectivity for the following two predicates as accurately as possible. Assume that values are evenly distributed inside the buckets.

i) **credits** ≤ 410

ii) **credits** > 141

b) You are given the following query in relational algebra, with relation sizes and *uncorrelated* selectivities as in the tables below. Assume that *S* and *T* only share the column name *a* which is a foreign key from *S* to *T*. *T* and *U* only share the column name *b* which is a foreign key from *T* to *U*. *S* and *U* do not share any column names.

$$\sigma_{\varphi} S \bowtie T \bowtie \sigma_{\psi} U$$

Estimate the number of tuples in the result and determine the join order that minimizes the size of the intermediate results. *Don't forget about the uniformity assumption!*

Relation	Size
S	400
T	10 000
U	20

Selection	Selectivity
σ_{φ}	0.25
σ_{ψ}	0.6

Question 5:

(4)

As so often, you are casually discussing database systems with some colleagues at a party. The night turns serious as one of your colleagues suddenly claims:

“If I join a very small relation R with a very large relation S , there are B^+ -tree indices on all join attributes and I expect that the result $R \bowtie S$ will be empty, then I should always use a sort-merge join.”

Do you agree? Briefly argue your position.

(You should assume that *very large* relations don't fit into the buffer, while *very small* relations always do.)

Overall: 20 points