

Group A

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

Exam 2 ON		SOLUTION KEY		29.05.2019
		Advanced Database Systems (184.780)		GROUP A
Matrikelnr.	Last Name		First Name	

Duration: 75 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. The optimal block size in HDFS is between 4 and 64 kilobytes.
2. Consider a table T in Hive for which neither partitioning nor bucketing is defined, and suppose that T is stored in HDFS. Then T is represented in HDFS as a directory with one or more files.
3. Consider a one-round MapReduce program without combiners for multiplying two $n \times n$ matrices with n^2 reducers, such that each reducer computes one entry of the resulting matrix. Then the program can be written in such a way that each reducer receives at most $2n$ values.
4. In Spark, several wide transformations can be grouped into one stage.

Question 2:

(4)

You are faced with the problem of implementing a join of two relations R and S in MapReduce in such a way that the communication cost is as low as possible. By surfing on the internet you have found an interesting web site, where it is stated that the semi-broadcast join is unbeatable. Please comment on this statement as follows:

- a) Explain the required MapReduce rounds of the semi-broadcast join and the communication cost of each round. [3 credits]

see slides 46/7 of the slide set II.1: MapReduce

b) Choose *one* alternative join method and explain how it works and what its communication cost is. [1 credit]

Question 3:

(4)

In the form of relational algebra implemented in SQL, relations are not sets, but bags; that is, tuples are allowed to appear more than once. Consequently, there are extended definitions of union, intersection, and difference for bags. Below we define the bag intersection. Sketch out a MapReduce algorithm for computing the following operation on bags R and S :

- **Bag Intersection**, defined to be the bag of tuples in which tuple t appears the minimum of the numbers of times it appears in R and S . Example: if t occurs 4 times in R and 2 times in S , then t should occur 2 times in the bag intersection of R and S .

In addition to this, also identify the communication cost of your algorithm, as a function of the input size.

Your Answer:

Question 4:

(4)

a) Evaluation of Spark/Scala [2 credits]

For this part of the question, you are given some Spark/Scala code and are supposed to evaluate the results. Don't worry about the "syntax" of your output, it is only important that it is clear that you understood the underlying semantics of Spark/Scala.

```
val fibs8 = List(0, 1, 1, 2, 3, 5, 8, 13)
val result = fibs8.filter(x => (x-1) % 3 == 0 && (x+2) % 2 == 1)
```

Note: '%' is the *modulo* operator (i.e., the rest after division).

```
val primes4 = List(2, 3, 5, 7)
def foo(num:Int) = num > 1 && ((2 to num-1 toList) forall (x => num % x != 0))
val result = primes4.flatMap(x => x :: (x to x+1 toList).filter(x => !foo(x)))
```

Note: 'x :: xs' prepends an element x to a list xs, thus producing a new list and '(x to y toList)' generates a list ranging from x to y (inclusively).

b) Writing Spark/Scala functions [2 credits]

You are given a list of random words, for instance:

```
val words = List("automaton", "language", "logic", "closure")
```

Produce a single list with every word duplicated (i.e. appearing twice), unless that word is shorter than 6 characters.

Note: You can use the "length" method to produce the length of any string. *You are expected to write a programmatic solution, writing down a static list of strings will get 0 credits.*

(4)

Note: underlines represent primary keys, *italics* represent foreign keys. We don't state the type of attributes here, just assume some reasonable defaults (string, int, float, ...)

Assume the dataset is loaded and the corresponding views have already been created in Spark SQL. For the dataframes, simply use the name of each table appended with “DF” (e.g.: hubDF, bulbDF, etc.). You don’t need to generate these dataframes, just assume they are already loaded and usable.

You are now given a number of queries, either in Spark SQL or in the Dataframe API. Your task will be to state an equivalent query of the other type, i.e. give a Spark SQL equivalent if the Dataframe API version is given, or vice-versa.

[illegible][illegible]

```
val query3DF = roomDF.join(hubDF, roomDF("lies_in") === hubDF("id"))
    .join(routinesDF, roomDF("id") === routinesDF("room"))
    .groupBy( roomDF("name"), hubDF("name"))
    .agg(count("room").as("numRoutines"))
    .select(hubDF("name").as("Hub"), roomDF("name").as("Room"), $"numRoutines")
    .orderBy(desc("numRoutines"))
```

```
val query4SQL = spark.sql("SELECT      b.name,b.model, AVG(c.degree)
                                FROM      bulb b, room r, hub h, compat c
                                WHERE      b.part_of = r.id AND r.lies_in = h.id AND c.bulb = b.id
                                           AND c.hub = h.id AND c.degree > '0.5'
                                           AND r.created_on > '2019-05-20'
                                           AND b.made_by != h.made_by
                                GROUP BY   b.name, b.model")
```