

## Group A

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

Exam 2 ON		<b>SOLUTION KEY</b>	07.10.2019
		Advanced Database Systems ( <b>184.780</b> )	<b>GROUP A</b>
Matrikelnr.	Last Name	First Name	

Duration: 75 minutes. Provide the solutions on the designated pages. **Good Luck!**

### Question 1:

(4 credits)

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  $\geq 0$  credits on the entire exam question 1.

1. In Spark, the filter operation may be applied to mutable variables but not to immutable variables.
2. In Spark, the user should combine several small computation steps to a single step in order to avoid performance loss.
3. In the MapReduce execution model, if a map worker fails, then all map tasks by this worker have to be redone.
4. In Hive, the user can choose between several execution engines. These choices include MapReduce and Spark.

**Question 2:**

(4 credits)

A colleague of yours, who is writing her/his first MapReduce programme, is asking you for advice: She/he has recently heard about combiners but did not understand what they are good for.

a) [2 credits] Briefly explain to her/him

- when is a combiner applicable,
- what is the goal of a combiner, and
- how is this goal achieved?

b) [2 credit] The application which your colleague had in mind involves the computation of the average function. Can a combiner be used for this purpose?

- If your answer is yes, give a brief sketch of the combiner.
- If your answer is no, give a short argument, why not?

Question 3:

(4 credits)

In this example you are asked to write a given type of join in MapReduce. The implementation details are left open, and there are multiple possible ways to implement this join in MapReduce. You will get full points as long as your solution sketches out a correct MapReduce program of the join type that is defined below.

- The **Left Outer Join** of  $R$  and  $S$  ( $R \bowtie S$ ), is the result of the (natural) join between  $R$  and  $S$ , combined with those tuples of  $R$  with no join partner, with the missing attributes from  $S$  replaced by NULLs:

<b>R</b>	A	B	C	<b>S</b>	B	C	D	<b>R <math>\bowtie</math> S</b>	A	B	C	D
	1	2	3		2	3	5		1	2	3	5
	2	2	4						2	2	4	NULL

Assume for your solution the given relational schema:  $R(A, B, C)$  and  $S(B, C, D)$ .

In addition to sketching out a MapReduce program, also identify the communication cost of your algorithm, as a function of the input size.

Your Answer:

**Question 4:**

(4 credits)

## 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 fact5 = List(1, 2, 6, 24, 120)
val result = fact5.filter(x => x % 2 == 0 && x % 3 == 0).map(x => x/6)
```

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

```
val numbers = List(2, 4, 4, 5, 7, 7, 8, 10, 14)
def foo(num:Int) = List(num, num*2, num+3)
val otherNumbers = List(2, 4, 7).flatMap(foo)
val result = numbers.sum - otherNumbers.sum
```

**Note:** '.sum' produces the sum of a list of integers

## b) Writing Spark/Scala functions [2 credits]

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

```
val words = List("madam", "letter", "rotator", "mammal")
```

A *palindrome* is defined as a word which is the same when read from left to right, or right to left. Examples: "racecar", "radar" or "level". For each word in your input, check if it's a palindrome, if not, add the warning ": no palindrome!" to the string. You can change the order of the list of words, but each word from the input must be present (as a prefix at least) in your output.

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

**Question 5:**

(4 credits)

You are given a database for a retailer.

**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, ...)

manufacturer (name, country, ceo, main\_product: *product.id*)  
product (id, name, first\_produced, made\_by: *manufacturer.name*)  
offers (id: *product.id*, brand: *retailer.brand*, price, created\_on)  
retailer (brand, country, no\_of\_stores)  
discount (by\_brand: *retailer.brand*, code, for\_product: *product.id*)

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.: productDF, discountDF, 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.

```
val query1DF = manufacturerDF.join(retailerDF, manufacturerDF("country") === retailerDF("country"))  
    .groupBy(manufacturerDF("name"))  
    .agg(count(retailerDF("brand")))
```

```
val query2DF = retailerDF.join(offerDF, offerDF("brand") === retailerDF("brand"))  
    .groupBy(retailerDF("brand"))  
    .agg(max("price").as("maxPrice"))  
    .select(retailerDF("brand"), $"maxPrice" + 100 )
```

```
val query3SQL = spark.sql("SELECT name, id
                            FROM manufacturer m, product p
                            WHERE p.made_by = m.name
                            AND NOT EXISTS (SELECT *
                                             FROM discount d
                                             WHERE d.for_product = p.id) ")
```

**Note:** You can use ".except" to form the difference of two dataframes **of the same schema**.

```
val query4SQL = spark.sql("SELECT    p.name,r.brand, AVG(o.price)
                            FROM      product p, retailer r, offers o
                            WHERE      p.id = o.id AND r.brand = o.brand
                                      AND r.no_of_stores > '4' AND o.price > '100.0'
                            GROUP BY  p.name, r.brand")
```

Overall: 20 points

Good luck!