

Exercise 1

Note: The following task can be solved as soon as Chapter 12 (Cardinality Estimation) from the script has been presented in the lecture. This will most likely be the case on May 13, 2019.

Assume you are given the following Relation R :

| R | |
|-----|-----|
| K | A |
| 0 | a |
| 1 | a |
| 2 | b |
| 3 | c |
| 4 | c |
| 5 | c |
| 6 | c |
| 7 | d |
| 8 | e |
| 9 | e |

Exercise 1 a)

Build a frequency vector for Attribute A of relation R .

Exercise 1 b)

The k^{th} frequency moment of a frequency vector f is defined as

$$F_k(f) = \sum_{i=1}^n f_i^k.$$

Compute the frequency moments for $k \in \{0, 1, 2\}$. What is the logical meaning of each of the values?

Exercise 1 c)

Sketches can be used to approximate the size of a join. Read *Sketches for Size of Join Estimation* by Rusu and Dobra, Section 1.1, and study Examples 2 (Section 3.1) and 3 (Section 3.2). In order to understand the examples, you may have to read more of Section (2 and) 3.

<http://faculty.ucmerced.edu/frusu/Papers/Journal/2008-05-tods.pdf>

Exercise 2

Note: The following task can be solved as soon as Chapter 13 (Parallelism) from the script has been presented in the lecture.

Amdahl's law allows you to compute the theoretical speed-up of the execution of a task when using n processors instead of 1. p denotes the proportion of the task that is parallelizable. The formula is given by

$$\frac{1}{(1-p) + \frac{p}{n}}$$

Exercise 2 a)

Compute the speed-up according to Amdahl's law for $p \in \{0.1, 0.5, 0.9\}$ and $n \in \{2, 8, 32\}$.

Exercise 2 b)

What are the implicit assumptions of Amdahl's law?

Exercise 2 c)

In the script in Chapter 13, p. 117, you can find a figure showing the speedup factor (y-axis) according to Amdahl's law depending on the fraction $p \in [0, 1]$ of parallelizable instructions (x-axis) for different number of processors n .

Sketch a diagram that shows the speedup factor on the y-axis and the number of processors n on the x-axis. Draw graphs for $p \in \{0.50, 0.75, 0.90, 0.95\}$. To which limits do the curves converge for large n ?

Exercise 2 d)

Discuss the difference between *inter-query parallelism* and *intra-query parallelism*.

How to Study for the DBS II Exam

Recall the topics covered in the lecture:

- Hardware
- Hashing
- Compression
- Storage Layout
- Physical Algebra: Processing Modes & Implementation
- Expression Evaluation
- Indexing
- Boolean Expressions
- Cardinality Estimation
- Parallelism

We spent a fair amount of time on Hardware and Hashing. However, we are neither electrical engineers nor mathematicians, so questions regarding these topics will either be not super deep or *from the perspective of a computer scientist*.

Proofs and product-specific stuff is of minor importance for the exam.

Programming is necessary to implement the presented concepts and will be important for the exam. Despite we used C++, this is a not a C++-lecture. In terms of the exam, basic programming concepts are important. In the exam, if you do not know how to answer a question using C++, use pseudo code. Correct C++ code will always give you full points. Pseudo code tends to be more vague and hide certain implementation details, e.g. regarding efficiency. Therefore, make sure to be as specific as possible, especially with regard to the data structures used. The storage layout chapter is an ideal candidate for programming questions. It is suggested to take a look at the slides/script and the exercises.

Expect the exam questions to check for an basic understanding of the presented topics. We also ask questions like *Name n properties of Y* and *Give the definition of X* since these are easily gained points for many students. A few (hard) questions that are designed to be *unexpected* are possible and play a role if you aim for very good grade.