

Database Systems II – Exercise #6

Sheet #6: Physical Algebra Implementation, BitSliceH

Daniel Flachs

Chair of Practical Computer Science III:
Database Management Systems

03/04/2019



- 1 Course Evaluation
- 2 Exercise Sheet #6
 - Task 2
 - Task 1

Contents

1 Course Evaluation

2 Exercise Sheet #6

- Task 2
- Task 1

Course Evaluation

Thank you very much for your feedback!

The written comments are always a great help, so please fill in the blank text areas at the back of the page (section 6).

Some inspiration:

- When you look at the exercise sheet, would you rather do a theoretical (pen and paper) or a programming exercise?
- Presentation: More slides? More blackboard? Something different?
- How could the presentation of coding solutions be improved?
- Should there be more time to discuss general questions regarding the lecture content, i. e., should the discussion of exercises be shorter?

Contents

1 Course Evaluation

2 Exercise Sheet #6

- Task 2
- Task 1

Task 2

You are given the following database relation:

| Persons | | |
|---------|---------|-----------|
| id | name | haircolor |
| 0 | Tom | brown |
| 1 | Olivia | blond |
| 2 | Esteban | black |
| 3 | Gaspar | black |
| 4 | Davy | brown |
| 5 | Jade | red |
| 6 | Daniel | brown |
| 7 | Clemens | blond |

- a) Assume that the value of the attribute `haircolor` is either black, brown, blond or red. The attribute is not nullable. How many bits do you need to encode these hair colors? Give a concrete encoding.

Solution: Two bits suffice to encode four different values. Let $k = 2$ denote the number of bits. The choice of the concrete encoding is arbitrary. One possibility is

black \mapsto 00

brown \mapsto 01

blond \mapsto 10

red \mapsto 11.

Task 2

b) Execute the following query by hand:

```
SELECT name
FROM Persons
WHERE haircolor
      <> 'blond';
```

| Persons | | | |
|---------|---------|-----------|----------|
| id | name | haircolor | encoding |
| 0 | Tom | brown | → 01 |
| 1 | Olivia | blond | → 10 |
| 2 | Esteban | black | → 00 |
| 3 | Gaspar | black | → 00 |
| 4 | Davy | brown | → 01 |
| 5 | Jade | red | → 11 |
| 6 | Daniel | brown | → 01 |
| 7 | Clemens | blond | → 10 |

Use the *BitSliceH* method presented in the lecture (Script, pp. 61 ff.) to retain only the hair colors that fulfill the selection predicate. Use a register size of $w = 16$. Show all intermediate steps. Indicate which tuples qualify.

Solution: see board.

Task 1

Download the zip archive from the website and try to make yourself familiar with the classes in `./physAlgEx6/PhysAlgebra.hh`. The classes implement the operators of a physical algebra with the following characteristics:

- push-based (memory flow and control flow)
- tuple-at-a-time processing
- processes tuples in row format
- assumes data is stored in row format.

Task 1a

Implement the `init()`, `step()` and `fin()` functions of the `Selection` operator of the physical algebra. If you get stuck, you may want to have a look at the solution code.

Task 1b

Implement the following queries:

- (i)

```
SELECT name
FROM test
WHERE name = 'Olivia';
```

- (ii)

```
SELECT *
FROM crimeInAtlanta2017
WHERE neighborhood = 'Peachtree_Hills';
```

- (iii)

```
SELECT location , MaxOfnum_victims
FROM crimeInAtlanta2017
WHERE x BETWEEN -84.36 AND -84.35
      AND y BETWEEN 33.73 AND 33.74;
```