



#### Query Tuning Query Processing

### Equivalence Rules – Examples

- Selection operations are commutative:  $\sigma_{\theta_1}(\sigma_{\theta_2}(E)) = \sigma_{\theta_2}(\sigma_{\theta_1}(E))$ 
  - E is a relation (table)
  - $\theta_1$  and  $\theta_2$  are conditions on attributes, e.g. *E.sallary* < 2500
  - $\sigma_{\theta}$  selects all tuples that satisfy  $\theta$
- Selection distributes over the theta-join operation if θ<sub>1</sub> involves only attributes of E<sub>1</sub> and θ<sub>2</sub> only attributes of E<sub>2</sub>:

 $\sigma_{\theta_1 \land \theta_2}(E_1 \bowtie_{\theta} E_2) = (\sigma_{\theta_1}(E_1)) \bowtie_{\theta} (\sigma_{\theta_2}(E_2))$ 

- $\bowtie_{\theta}$  is the theta-join; it pairs tuples from the input relations (e.g.,  $E_1$  and  $E_2$ ) that satisfy condition  $\theta$ , e.g.  $E_1$ .account $ID = E_2.ID$
- Natural join is associative:  $(E_1 \bowtie E_2) \bowtie E_3 = E_1 \bowtie (E_2 \bowtie E_3)$ 
  - the join condition in the natural join is equality on all attributes of the two input relations that have the same name
- Many other rules can be found in Silberschatz et al., "Database System Concepts"

DBT - Query Tuning

Query Tuning Query Processing Equivalence Rules – Example Query

• Equivalent relational algebra expressions:

 $\begin{array}{c} \Pi_{customer-name} \\ \square \\ \sigma_{branch-city = Brooklyn} \\ \land balance < 1000 \\ \square \\ branch \\ depositor \end{array} = \\ \begin{array}{c} M \\ \sigma_{branch-city = Brooklyn} \\ \sigma_{branch-city = Brooklyn} \\ \sigma_{branch-city = Brooklyn} \\ \sigma_{balance < 1000} \\ \square \\ branch \\ account \end{array}$ 

### Query Tuning Query Processing

### Equivalence Rules – Example Query

Schema:

branch(branch-name, branch-city, assets)
account(account-number, branch-name, balance)
depositor(customer-name,account-number)

• Query:

SELECT customer-name FROM branch, account, depositor WHERE branch-city=Brooklyn AND balance < 1000 AND branch.branch-name = account.branch-name AND account.account-number = depositor.account-number

#### Nikolaus Augsten (DIS)

DBT – Query Tuning
Query Tuning Query Processing

#### Unit 2 – WS 2016/17 10 / 58

## B) Annotation: Creating Query Plans

- Algebra expression is not a query plan.
- Additional decisions required:
  - which indexes to use, for example, for joins and selects?
  - which algorithms to use, for example, sort-merge vs. hash join?
  - materialize intermediate results or pipeline them?
  - etc.
- Each relational algebra expression can result in many query plans.
- Some query plans may be better than others!

Nikolaus Augsten (DIS)

Unit 2 - WS 2016/17

9 / 58



- selection size estimation
- join size estimation
- projection size estimation
- Problems:
  - cost can only be estimated
  - updating statistics is expensive, thus they are often out of date

Unit 2 – WS 2016/17 15 / 58 Nikolaus Augsten (DIS)

• Examples of heuristics:

• perform projections early • avoid Cartesian products

• perform selections as early as possible

| Query Tuning Query Processing  | Query Tuning Query Processing  |  |  |
|--|--|--|--|
| 3. Execution Engine  | Query Tuning and Query Optimization  |  |  |
| The execution engine<br>• receives query plan from optimizer<br>• executes plan and returns query result to user   | <ul> <li>Optimizers are not perfect:</li> <li>transformations produce only a subset of all possible query plans</li> <li>only a subset of possible annotations might be considered</li> <li>cost of query plans can only be estimated</li> <li>Query Tuning: Make life easier for your query optimizer!</li> </ul> |  |  |
| Nikolaus Augsten (DIS) DBT – Query Tuning Unit 2 – WS 2016/17 17 / 58<br>Query Tuning Problematic Queries<br>Outline                                       | Nikolaus Augsten (DI5) DBT – Query Tuning Unit 2 – WS 2016/17 18 / 58 Query Tuning Problematic Queries Which Queries Should Be Rewritten?  |  |  |
| <ol> <li>Query Tuning</li> <li>Query Processing</li> <li>Problematic Queries</li> <li>Minimizing DISTINCTs</li> <li>Rewriting of Nested Queries</li> </ol> | <ul> <li>Rewrite queries that run "too slow"</li> <li>How to find these queries? <ul> <li>query issues far too many disc accesses,</li> <li>for example, point query scans an entire table</li> <li>you look at the query plan and see that relevant indexes are not used</li> </ul> </li> </ul>                   |  |  |
| Nikolaus Augsten (DIS) DBT – Query Tuning Unit 2 – WS 2016/17 19 / 58  | Nikolaus Augsten (DIS) DBT – Query Tuning Unit 2 – WS 2016/17 20 / 58  |  |  |

#### Query Tuning Problematic Queries

### Running Example

- Employee(<u>ssnum,name</u>,manager,dept,salary,numfriends)
  - clustering index on ssnum
  - non-clustering index on name
  - non-clustering index on dept
  - keys: ssnum, name
- Students(<u>ssnum</u>,<u>name</u>,course,grade)
  - clustering index on ssnum
  - non-clustering index on name
  - keys: ssnum, name
- Techdept(dept,manager,location)
  - clustering index on dept
  - key: dept
  - manager may manage many departments
  - a location may contain many departments

#### Nikolaus Augsten (DIS)

DBT - Query Tuning

Query Tuning Problematic Queries

### Non-Correlated Subqueries

- Many systems handle subqueries inefficiently.
- Non-correlated: attributes of outer query not used in inner query.
- Query:

SELECT ssnum FROM Employee

WHERE dept IN (SELECT dept FROM Techdept)

- May lead to inefficient evaluation:
  - check for each employee whether they are in Techdept

DBT – Query Tuning

- index on Employee.dept not used!
- Equivalent query:

### SELECT ssnum FROM Employee, Techdept

WHERE Employee.dept = Techdept.dept

- Efficient evaluation:
  - look up employees for each dept in Techdept
  - use index on Employee.dept

Unit 2 - WS 2016/17

21 / 58

### Query Tuning Problematic Queries

DISTINCT

- How can DISTINCT hurt?
  - DISTINCT forces sort or other overhead.
  - If not necessary, it should be avoided.
- Query: Find employees who work in the information systems department.

SELECT DISTINCT ssnum FROM Employee WHERE dept = 'information systems'

- DISTINCT not necessary:
  - ssnum is a key of Employee, so it is also a key of a subset of Employee.
  - Note: Since an index is defined on ssnum, there is likely to be no overhead in this particular examples.

#### Nikolaus Augsten (DIS)

Query Tuning Problematic Queries

DBT - Query Tuning

### Temporary Tables

- Temporary tables can hurt in the following ways:
  - force operations to be performed in suboptimal order (optimizer often does a very good job!)
  - creating temporary tables i.s.s.<sup>1</sup> causes catalog update possible concurrency control bottleneck
  - system may miss opportunity to use index
- Temporary tables are good:
  - to rewrite complicated correlated subqueries
  - to avoid ORDER BYs and scans in specific cases (see example)

<sup>1</sup>in some systems Nikolaus Augsten (DIS) Unit 2 - WS 2016/17

22 / 58

#### Query Tuning Problematic Queries

### Unnecessary Temporary Table

• Query: Find all IT department employees who earn more than 40000. SELECT \* INTO Temp

FROM Employee WHERE salary > 40000

SELECT ssnum FROM Temp WHERE Temp.dept = 'IT'

- Inefficient SQL:
  - index on dept can not be used
  - overhead to create Temp table (materialization vs. pipelining)
- Efficient SQL:

SELECT ssnum FROM Employee WHERE Employee.dept = 'IT' AND salary > 40000

Nikolaus Augsten (DIS)

DBT – Query Tuning

Query Tuning Problematic Queries

### Don't use HAVING where WHERE is enough

- Query: Find average salary of the IT department.
- Inefficient SQL:

SELECT AVG(salary) as avgsalary, dept FROM Employee GROUP BY dept HAVING dept = 'IT'

- Problem: May first compute average for employees of all departments.
- Efficient SQL: Compute average only for relevant employees.

```
SELECT AVG(salary) as avgsalary, dept
FROM Employee
WHERE dept = 'IT'
GROUP BY dept
```

### Joins: Use Clustering Indexes and Numeric Values

- Query: Find all students who are also employees.
- Inefficient SQL:

SELECT Employee.ssnum FROM Employee, Student WHERE Employee.name = Student.name

• Efficient SQL:

SELECT Employee.ssnum FROM Employee, Student WHERE Employee.ssnum = Student.ssnum

- Benefits:
  - Join on two clustering indexes allows merge join (fast!).
  - Numerical equality is faster evaluated than string equality.

#### Nikolaus Augsten (DIS)

DBT – Query Tuning

```
Query Tuning Problematic Queries
```

```
Use Views with Care (I/II)
```

- Views: macros for queries
  - queries look simpler
  - but are never faster and sometimes slower
- Creating a view:

CREATE VIEW Techlocation AS SELECT ssnum, Techdept.dept, location

FROM Employee, Techdept

- WHERE Employee.dept = Techdept.dept
- Using the view:

SELECT location FROM Techlocation WHERE ssnum = 452354786

• System expands view and executes:

SELECT location FROM Employee, Techdept WHERE Employee.dept = Techdept.dept AND ssnum = 452354786

Unit 2 – WS 2016/17 27 / 58

Unit 2 – WS 2016/17 25 / 58

Unit 2 - WS 2016/17 26 / 58

#### Query Tuning Problematic Queries

### Use Views with Care (II/II)

- Query: Get the department name for the employee with social security number 452354786 (who works in a technical department).
- Example of an inefficient SQL:

SELECT dept FROM Techlocation WHERE ssnum = 452354786

• This SQL expands to:

SELECT Techdept.dept
FROM Employee, Techdept
WHERE Employee.dept = Techdept.dept
AND ssnum = 452354786

• But there is a more efficient SQL (no join!) doing the same thing:

DBT - Query Tuning

SELECT dept FROM Employee WHERE ssnum = 452354786

Nikolaus Augsten (DIS)

Unit 2 – WS 2016/17

### Query Tuning Problematic Queries System Peculiarity: Order in FROM clause

- Order in FROM clause should be irrelevant.
- However: For long joins (e.g., more than 8 tables) and in some systems the order matters.
- How to figure out? Check query plan!

### System Peculiarity: Indexes and OR

- Some systems never use indexes when conditions are OR-connected.
- Query: Find employees with name Smith or who are in the acquisitions department.

SELECT Employee.ssnum FROM Employee WHERE Employee.name = 'Smith' OR Employee.dept = 'acquisitions'

Fix: use UNION instead of OR

SELECT Employee.ssnum FROM Employee WHERE Employee.name = 'Smith' UNION SELECT Employee.ssnum FROM Employee WHERE Employee.dept = 'acquisitions'

Nikolaus Augsten (DIS)

DBT – Query Tuning

Unit 2 - WS 2016/17 30 / 58

#### Query Tuning Problematic Queries

### Experimental Evaluation

### Throughput increase in percent.



Unit 2 - WS 2016/17 31 / 58

29 / 58





#### Query Tuning Minimizing DISTINCTs

### No-Duplicate Guarantee – Examples

- This query also returns no duplicates: SELECT ssnum, Techdept.dept FROM Employee, Techdept
- Reason: as before,

Nikolaus Augsten (DIS)

• both Techdept and Employee are privileged table

DBT - Query Tuning

# Query Tuning Minimizing DISTINCTS No-Duplicate Guarantee – Examples

- This query returns duplicates: (even if Student.name is a key)
   SELECT Student.ssnum
   FROM Student, Employee, Techdept
   WHERE Student.name = Employee.name
   AND Employee.manager = Techdept.manager
- Reason:
  - join attribute Techdept.manager is not key
  - thus Techdept does not reach Employee (and Student)

#### Query Tuning Minimizing DISTINCTs

### No-Duplicate Guarantee – Examples

• This query returns no duplicates: (even if Student.name is not a key)

SELECT Student.ssnum FROM Student, Employee, Techdept WHERE Student.name = Employee.name AND Employee.dept = Techdept.dept

- Reason:
  - join attribute Employee.name is a key, thus Employee reaches privileged table Student
  - join attribute Techdept.dept is a key thus Techdept reaches Employee
  - transitivity: Techdept reaches Employee and Employee reaches Student, thus Techdept reaches Student

#### Nikolaus Augsten (DIS)

DBT – Query Tuning

### Query Tuning Minimizing DISTINCTs

### No-Duplicate Guarantee – Examples

• Try the example queries on the following instance (keys underlined):

| • Employee( <u>ssnum</u> , <u>name</u> , manager, dept) |       |         |             |  |
|---|-------|---------|-------------|--|
| ssnum   | name  | manager | dept        |  |
| 1   | Peter | John    | IT          |  |
| 2   | Rose  | Mary    | Development |  |
| • Techdept(dept, manager)                               |       |         |             |  |
| dept  | dept  |         |             |  |
| IT  |       | John    |             |  |
| Development   |       | Mary    |             |  |
| Production  |       | John    |             |  |
| • Students( <u>ssnum</u> , name)                        |       |         |             |  |
| ssnum   | name  |         |             |  |
| 5   | Peter |         |             |  |
| 6   | Peter |         |             |  |
|   |       |         |             |  |
|   |       |         |             |  |

Unit 2 – WS 2016/17

41 / 58

Unit 2 - WS 2016/17 42 / 58





Unit 2 – WS 2016/17 51 / 58

Nikolaus Augsten (DIS)



#### Query Tuning Rewriting of Nested Queries

### The Count Bug



### The Count Bug

Nikolaus Augsten (DIS)

- Consider for example an employee Jane:
  - Jane is not in a technical department (Techdept).
  - Jane has no friends (Employee.numfriends = 0)
- Original (nested) query:
  - since Jane is not in a technical department, inner query is empty
  - but  $COUNT(\emptyset)=0$ , thus Jane is in the result set!
- Rewritten query with temporary table:
  - Jane not in a technical department and does not survive the join

DBT - Query Tuning

Unit 2 - WS 2016/17

58 / 58

• thus Jane is not in the result set