

An Overview of Query Optimization in Relational Systems

Surajit Chaudhuri
Microsoft Research
surajitc@microsoft.com
<http://research.microsoft.com/~surajitc>

Outline

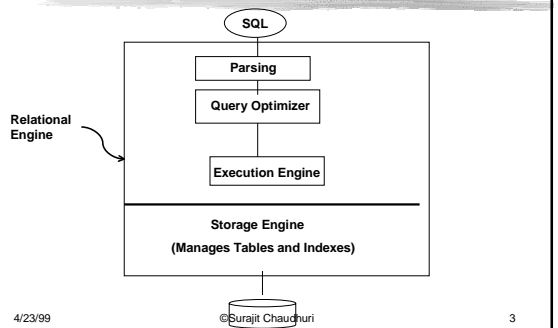
- Preliminaries
 - ▮ Relational query engine
 - ▮ Operator Trees
- Query Optimization Framework
- Building Blocks
- Tuning Optimizers
- Active Areas of Research
- Conclusion

4/23/99

©Surajit Chaudhuri

2

Relational DBMS Components



4/23/99

©Surajit Chaudhuri

3

Scan and Selection Operators

- Scan([index], table, predicate)
 - ▮ Sequential Scan
 - ▮ Indexscan: Which index(es) to use?
 - ▮ Always push down "index-evaluable" predicates
- Filter(table, predicate)

4/23/99

©Surajit Chaudhuri

4

Join Operators

- Join(Merge, R1, R2, R1.a = R2.a)
 - ▮ Requires sorted order on R1.a, R2.a
 - ▮ Output is a sorted order
 - ▮ Pipelined
- Join(Nested-Loop, R1, R2, R1.a = R2.b)
 - ▮ Sorted inputs of no consequence
 - ▮ Output has the same sort order as R1.a
 - ▮ Pipelined

4/23/99

©Surajit Chaudhuri

5

"Generic" Join Operators

- Join([method], outer, inner, join-predicate)
 - ▮ Asymmetric
 - ▮ Effect of physical properties of input streams (e.g., sorted input)
 - ▮ Physical properties of output stream (e.g., sorted)
 - ▮ Pipelined v.s. Blocking (Nested Loop v.s. Sort-Merge)

4/23/99

©Surajit Chaudhuri

6

Sort & Hash

- Sort($R_1, \{A,B,C\}$, parameters)
 - ▮ Blocking (only in part)
- Build-Hash($R_1, \{C,D\}$, parameters)
 - ▮ Blocking
- Probe-Hash($R_2, \{C,D\}$, hash-table)
 - ▮ Pipelining

4/23/99

©Surajit Chaudhuri

7

Generic View of RE Operators


- Input: One or more data streams
- Output: One data stream
- Implementation
 - ▮ open()
 - ▮ getNext()
 - ▮ close()
- Pipelined/Blocking

4/23/99

©Surajit Chaudhuri

8

RE Operator Tree



4/23/99

©Surajit Chaudhuri

9

Execution of an Operator Tree

- Demand-driven architecture is the simplest
- open() is propagated from the root
- getNext() at the root is propagated
- If getNext() at the root fails to return a new tuple, then no more answers for the query

4/23/99

©Surajit Chaudhuri

10

Outline

- Preliminaries
- Query Optimization Framework
 - ▮ Why do it?
 - ▮ What information needs to be considered?
- Building Blocks
- Tuning Optimizers
- Active Areas of Research
- Conclusion

4/23/99

©Surajit Chaudhuri

11

Query Optimization

- Compile a SQL query in an operator tree
- Approach 1:
 - ▮ Take a SQL expression
 - ▮ Represent as a tree
 - ▮ Turn algebraic operators into RE operators
- Why is this not adequate?

4/23/99

©Surajit Chaudhuri

12

Observations

- Selections always reduce relation size
- Index scan may/may not be more efficient
- Order of evaluation of joins irrelevant for correctness
 - ┆ But, order determines efficiency
- Total sales for Products for Q298 from NW
 - ┆ Sum sales of all products, join with Products
 - ┆ Join products with Sale, then sum

4/23/99

©Surajit Chaudhuri

13

Richness of Choice

- Algebraic properties allow semantically equivalent algebraic trees for a query
- Multiple implementation techniques for each algebraic operator
- Costs of the alternatives may be widely different depending on statistical properties of data

4/23/99

©Surajit Chaudhuri

14

Goal of Query Optimizer

- Compile a SQL query in an operator tree
- Find the tree with least cost subject to
 - ┆ Equivalence transformations
 - ┆ Available RE operators
 - ┆ Database statistics
- Explore the space of trees efficiently
- Optimizer helps achieve
 - ┆ Data Independence

4/23/99

©Surajit Chaudhuri

15

Implementing an Optimizer

- Program Optimization
- Search Algorithms/Planning
- Combinatorial Optimization
- We will revisit this issue in the next lecture..

4/23/99

©Surajit Chaudhuri

16

A Framework for Query Optimization

- Equivalence Transformations
 - ┆ Algebraic properties
 - ┆ Implementation options
- Estimation Model
 - ┆ Needs to estimate cost of an operator tree (incrementally)
- Tree-Finder ("search")
 - ┆ Fast, Memory-efficient

4/23/99

©Surajit Chaudhuri

17

Outline

- Preliminaries
- Query Optimization Framework
- Building Blocks
 - ┆ *Equivalence Transformations*
 - ┆ Statistical Model
 - ┆ Tree-Finder
- Tuning Optimizers
- Active Areas of Research
- Conclusion

4/23/99

©Surajit Chaudhuri

18

Examples of Transformations

- Join Ordering
- Commuting join and outer-join
- Commuting group by and join
- Optimize multi-block queries
 - ┆ Collapse multi-block query to a single block query
 - ┆ Optimize across multiple query blocks

4/23/99

©Surajit Chaudhuri

19

ASPJ Queries

```
Select A.a, Sum(B.z)
From A, B, C
Where A.x = B.x and B.y = C.y
Group By A.a
Order By A.a
```

4/23/99

©Surajit Chaudhuri

20

Implementation Transformations

- Scan
 - ┆ B+ tree index scan with (*sargable*) Predicate: Between and its degenerate forms
 - ┆ Sequential scan
- Filter
 - ┆ Any Boolean expression
- Join
 - ┆ Sort-Merge, Nested-loop, Indexed Nested-loop

4/23/99

©Surajit Chaudhuri

21

Join Ordering

- Select and Join commute
 - ┆ $\text{Filter}(\text{Join}(A,B), a) = \text{Join}(\text{Filter}(A,a), B)$
- Joins are associative and commutative :
 - ┆ $\text{Join}(\text{Join}(A,B), C) = \text{Join}(\text{Join}(B,A), C)$
 - ┆ $\text{Join}(\text{Join}(A,C), B) = \text{Join}(\text{Join}(A,B), C)$
 - ┆ Many equivalent expressions (How many?)
- How about n-ary joins?


4/23/99

©Surajit Chaudhuri

22

Transformation => Space

- Shape of join trees
 - ┆ Restricted use of AC properties
- Linear Join Trees
- Bushy Join Trees



4/23/99

©Surajit Chaudhuri

23

Join and OuterJoin

- Do outer-join operators commute?
- Example:
 - ┆ $(R \text{ LOJ } S)$: Preserves R
 - ┆ $\text{Join}(R, S \text{ LOJ } T) = ?$
- Goal of Transformations: Isolate blocks of natural join

4/23/99

©Surajit Chaudhuri

24

Group By and Join

- Select .. From .. Where .. Group By ..
 - ┆ Traditionally, execution of group-by follows execution of joins
- Total sales for Products for Q298 from NW
 - ┆ Sum sales of all products, join with Products
 - ┆ Join products with Sale, then sum

4/23/99


©Surajit Chaudhuri

25

Operator Trees: Group By and Join

- Schema:
 - ┆ Product(pid, unitprice, ..)
 - ┆ Sales(tid, date, store, pid, units)

Trees:



4/23/99

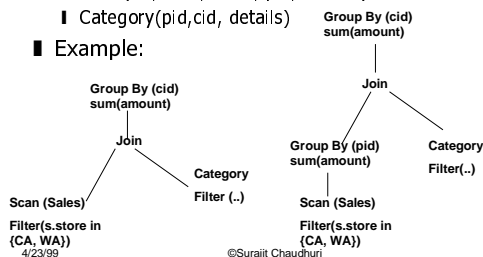
©Surajit Chaudhuri

26

Introducing Operator: Group By and Join

- Schema:
 - ┆ Sales(tid, date, store, pid, amount)
 - ┆ Category(pid, cid, details)

Example:



4/23/99

©Surajit Chaudhuri

27

Prerequisites: Group By and Join

- Schema constraints, arbitrary aggregation functions
- No schema constraints, but properties of aggregate functions
 - ┆ $Agg(S1 \cup S2) = f(Agg(S1), Agg(S2))$
 - ┆ May sometime require use of derived columns

4/23/99

©Surajit Chaudhuri

28

Group By and Join: Pros and Cons

- (+) "Pushing down" group by past a join:
 - ┆ Group By "collapses" an equivalence class
 - ┆ Therefore, may reduce cost of subsequent joins
 - ┆ Can be pipelined with index scans
- (-) Application needs to be cost based since
 - ┆ The cost of group by may be increased
 - ┆ Access methods on base tables may no longer be useful for the join
- Encapsulating a "transformation" is not easy

4/23/99

©Surajit Chaudhuri

29

Multi-Block Queries

- Multi-block structure arises due to
 - ┆ views with aggregates
 - ┆ table expressions
 - ┆ nested sub-queries
- Techniques for Optimization
 - ┆ Merge into a single block
 - ┆ Share information across blocks

4/23/99

©Surajit Chaudhuri

30