

CMU SCS

Carnegie Mellon Univ.
Dept. of Computer Science
15-415/615 - DB Applications

C. Faloutsos – A. Pavlo
Lecture#15: Query Optimization

CMU SCS

Administrivia

- HW5 is due **Wed Oct 28th**.
- Mid-terms are available for **viewing**
 - Marilyn Walgora's office (GHC 8120)
 - 9:00am-11:45am and 1:00pm-4:15pm
 - Bring your CMU id

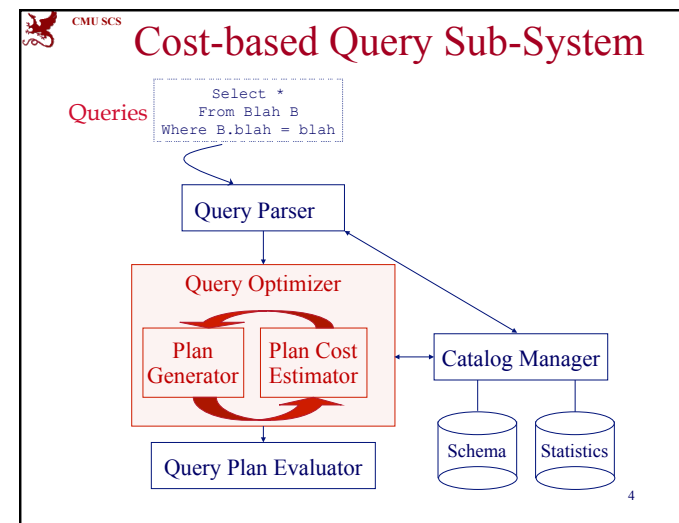
Faloutsos/Pavlo CMU SCS 15-415/615 2

CMU SCS

Today's Class

- History & Background
- Relational Algebra Equivalences
- Plan Cost Estimation
- Plan Enumeration

Faloutsos/Pavlo CMU SCS 15-415/615 3



CMU SCS

Query Optimization


- Remember that SQL is declarative.
 - User tells the DBMS *what* answer they want, not *how* to get the answer.
- There can be a big difference in performance based on plan is used:
 - See last week: **5.7 days** vs. **45 seconds**

Faloutsos/Pavlo CMU SCS 15-415/615 5

CMU SCS

1970s – Relational Model

- Ted Codd saw the maintenance overhead for IMS/Codasyl.
- Proposed database abstraction based on relations:
 - Store database in simple data structures.
 - Access it through high-level language.
 - Physical storage left up to implementation.



Codd

Faloutsos/Pavlo CMU SCS 15-415/615 6

CMU SCS

IBM System R

- Skunkworks project at IBM Research in San Jose to implement Codd's ideas.
- Had to figure out all of the things that we are discussing in this course themselves.
- IBM never commercialized **System R**.

Faloutsos/Pavlo CMU SCS 15-415/615 7

CMU SCS

IBM System R

- First implementation of a query optimizer.
- People argued that the DBMS could never choose a query plan better than what a human could write.
- A lot of the concepts from System R's optimizer are still used today.

Faloutsos/Pavlo CMU SCS 15-415/615 8

CMU SCS

Today's Class

- History & Background
- Relational Algebra Equivalences
- Plan Cost Estimation
- Plan Enumeration
- Nested Sub-queries

Faloutsos/Pavlo CMU SCS 15-415/615 9

CMU SCS

Relational Algebra Equivalences

- A query can be expressed in different ways.
- The optimizer considers variations and choose the one with the lowest cost.
- How do we know whether two queries are equivalent?

Faloutsos/Pavlo CMU SCS 15-415/615 10

CMU SCS

Relational Algebra Equivalences

- Two relational algebra expressions are *equivalent* if they generate the same set of tuples.

Faloutsos/Pavlo CMU SCS 15-415/615 11

CMU SCS

Predicate Pushdown

SELECT cname, amt
FROM customer, account
WHERE customer.acctno =
 account.acctno
AND account.amt > 1000

π cname, amt

σ amt>1000

\bowtie acctno=acctno

CUSTOMER ACCOUNT

➔

π cname, amt

\bowtie acctno=acctno

σ amt>1000

CUSTOMER ACCOUNT

Faloutsos/Pavlo CMU SCS 15-415/615 12

CMU SCS

Relational Algebra Equivalences

```
SELECT cname, amt
FROM customer, account
WHERE customer.acctno =
      account.acctno
AND account.amt > 1000
```

$$\pi_{\text{cname, amt}}(\sigma_{\text{amt}>1000}(\text{customer} \bowtie \text{account}))$$

$$=$$

$$\pi_{\text{cname, amt}}(\text{customer} \bowtie (\sigma_{\text{amt}>1000}(\text{account})))$$

Faloutsos/Pavlo CMU SCS 15-415/615 13

CMU SCS

Relational Algebra Equivalences

- **Selections:**
 - Perform them early
 - Break a complex predicate, and push

$$\sigma_{p1 \wedge p2 \wedge \dots \wedge pn}(R) = \sigma_{p1}(\sigma_{p2}(\dots \sigma_{pn}(R)) \dots)$$

- Simplify a complex predicate
 - $(X=Y \text{ AND } Y=3) \rightarrow X=3 \text{ AND } Y=3$

Faloutsos/Pavlo CMU SCS 15-415/615 14

CMU SCS

Relational Algebra Equivalences

- **Projections:**
 - Perform them early
 - Smaller tuples
 - Fewer tuples (if duplicates are eliminated)
 - Project out all attributes except the ones requested or required (e.g., joining attr.)

Faloutsos/Pavlo CMU SCS 15-415/615 15

CMU SCS

Projection Pushdown


```
SELECT cname, amt
FROM customer, account
WHERE customer.acctno =
      account.acctno
AND account.amt > 1000
```

$\pi_{\text{cname, amt}}$
 \downarrow
 $\sigma_{\text{amt}>1000}$
 \downarrow
 $\bowtie_{\text{acctno=acctno}}$
 $\swarrow \searrow$
CUSTOMER ACCOUNT

➔

$\pi_{\text{cname, amt}}$
 \downarrow
 $\bowtie_{\text{acctno=acctno}}$
 $\swarrow \searrow$
 $\sigma_{\text{amt}>1000}$
 \downarrow
 $\pi_{\text{acctno, amt}}$
 $\swarrow \searrow$
CUSTOMER ACCOUNT

Faloutsos/Pavlo CMU SCS 15-415/615 16




Relational Algebra Equivalences

- **Joins:**
 - Commutative, associative
$$R \bowtie S = S \bowtie R$$

$$(R \bowtie S) \bowtie T = R \bowtie (S \bowtie T)$$
- Q: How many different orderings are there for an n -way join?


Faloutsos/Pavlo CMU SCS 15-415/615 17



Relational Algebra Equivalences

- **Joins:** How many different orderings are there for an n -way join?
- A: [Catalan number](#) $\sim 4^n$
 - Exhaustive enumeration: too slow.
- We'll see in a second how an optimizer limits the search space...


Faloutsos/Pavlo CMU SCS 15-415/615 18



Today's Class

- History & Background
- Relational Algebra Equivalences
- Plan Cost Estimation
- Plan Enumeration

Faloutsos/Pavlo CMU SCS 15-415/615 19



Cost Estimation

- How long will a query take?
 - **CPU:** Small cost; tough to estimate
 - **Disk:** # of block transfers
 - **Memory:** Amount of DRAM used
 - **Network:** # of messages
- How many tuples will qualify?
- What statistics do we need to keep?

Faloutsos/Pavlo CMU SCS 15-415/615 20

CMU SCS

Cost Estimation – Statistics

- For each relation **R** we keep:
 - N_R → # tuples
 - S_R → size of tuple in bytes
 - $V(A,R)$ → # of distinct values of attribute 'A'

Faloutsos/Pavlo CMU SCS 15-415/615 21

CMU SCS

Derivable Statistics

- F_R → max# records/block
- B_R → # blocks
- $SC(A,R)$ → selection cardinality
avg# of records with A=given

Faloutsos/Pavlo CMU SCS 15-415/615 22

CMU SCS

Derivable Statistics

- $SC(A,R)$ → Selection Cardinality
avg# of records with A=given
→ $N_R / V(A,R)$
- Note that this assumes data uniformity
 - 10,000 students, 10 colleges – how many students in SCS?

Faloutsos/Pavlo CMU SCS 15-415/615 23

CMU SCS

Additional Statistics

- For index **i**:
 - F_i → average fanout (~50-100)
 - HT_i → # levels of index **i** (~2-3)
~ $\log(\#entries) / \log(F_i)$
 - LB_i → blocks at leaf level

Faloutsos/Pavlo CMU SCS 15-415/615 24

CMU SCS

Statistics

- Where do we store them?
- How often do we update them?
- Manual invocations:
 - Postgres: **ANALYZE**
 - MySQL: **ANALYZE TABLE**

Faloutsos/Pavlo CMU SCS 15-415/615 25

CMU SCS

Selection Statistics

- We saw simple predicates (**name="Trump"**)
- How about more complex predicates, like
 - **salary > 10000**
 - **age=30 AND jobTitle="Costermonger"**
- What is their selectivity?

Faloutsos/Pavlo CMU SCS 15-415/615 26

CMU SCS

Selections – Complex Predicates

- Selectivity **sel(P)** of predicate **P**:
 == fraction of tuples that qualify

$$\text{sel(P)} = \text{SC(P)} / N_R$$

Selection Cardinality

/

N_R

=

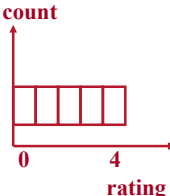
of tuples

Faloutsos/Pavlo CMU SCS 15-415/615 27

CMU SCS

Selections – Complex Predicates

- Assume that **V(rating, SAILORS)** has 5 distinct values (i.e., 0 to 4).
- simple predicate **P: A=constant**
 - $\text{sel(A=constant)} = 1/V(A,R)$
 - eg., $\text{sel(rating='2')} = 1/5$
- What if **V(A,R)** is unknown??



Faloutsos/Pavlo CMU SCS 15-415/615 28

CMU SCS

Selections – Complex Predicates

- Range Query: $\text{sel}(\text{rating} \geq '2')$
- $\text{sel}(A > a) = (A_{\max} - a) / (A_{\max} - A_{\min})$

Faloutsos/Pavlo CMU SCS 15-415/615 29

CMU SCS

Selections – Complex Predicates

- Negation: $\text{sel}(\text{rating} \neq '2')$
– $\text{sel}(\text{not } P) = 1 - \text{sel}(P)$
- Observation: selectivity \approx probability

Faloutsos/Pavlo CMU SCS 15-415/615 30

CMU SCS

Selections – Complex Predicates

- **Conjunction:**
 - $\text{sel}(\text{rating} = '2' \text{ and name LIKE 'C\%')$
 - $\text{sel}(P1 \wedge P2) = \text{sel}(P1) \cdot \text{sel}(P2)$
 - INDEPENDENCE ASSUMPTION

Faloutsos/Pavlo CMU SCS 15-415/615 31

CMU SCS

Selections – Complex Predicates

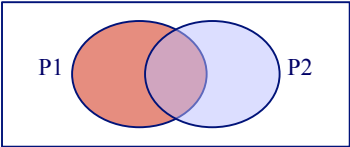
- **Disjunction:**
 - $\text{sel}(\text{rating} = '2' \text{ or name LIKE 'C\%')$
 - $\text{sel}(P1 \vee P2)$
= $\text{sel}(P1) + \text{sel}(P2) - \text{sel}(P1 \wedge P2)$
= $\text{sel}(P1) + \text{sel}(P2) - \text{sel}(P1) \cdot \text{sel}(P2)$
 - INDEPENDENCE ASSUMPTION, again

Faloutsos/Pavlo CMU SCS 15-415/615 32

CMU SCS

Selections – Complex Predicates

- **Disjunction, in general:**
 - $\text{sel}(P1 \text{ or } P2 \text{ or } \dots Pn) =$
 - $1 - (1 - \text{sel}(P1)) \cdot (1 - \text{sel}(P2)) \cdot \dots (1 - \text{sel}(Pn))$



Faloutsos/Pavlo CMU SCS 15-415/615 33

CMU SCS

Joins

- Q: Given a join of **R** and **S**, what is the range of possible result sizes in #of tuples?

Faloutsos/Pavlo CMU SCS 15-415/615 34

CMU SCS

Result Size Estimation for Joins

- General case: $R_{\text{cols}} \cap S_{\text{cols}} = \{A\}$ where A is not a key for either table.
- *Hint: for a given tuple of R, how many tuples of S will it match?*

Faloutsos/Pavlo CMU SCS 15-415/615 35

CMU SCS

Result Size Estimation for Joins

- General case: $R_{\text{cols}} \cap S_{\text{cols}} = \{A\}$ where A is not a key for either table.
 - Match each **R**-tuple with **S**-tuples:
 $\text{estSize} \approx N_R \cdot N_S / V(A,S)$
 - Symmetrically, for **S**:
 $\text{estSize} \approx N_R \cdot N_S / V(A,R)$
- Overall:
 - $\text{estSize} \approx N_R \cdot N_S / \max(\{V(A,S), V(A,R)\})$

Faloutsos/Pavlo CMU SCS 15-415/615 36

CMU SCS

Cost Estimations

- Our formulas are nice but we assume that data values are uniformly distributed.

of occurrences

Uniform Approximation of D

Distinct values of attribute

37

CMU SCS

Cost Estimations

- Our formulas are nice but we assume that data values are uniformly distributed.

Non-Uniform Distribution D

Bucket 1 Bucket 2 Bucket 3 Bucket 4 Bucket 5
 Count=8 Count=4 Count=15 Count=3 Count=15

38

CMU SCS

Histograms w/ Quantiles

- Allows the DBMS to have leverage better statistics about the data.

Equiwidth Histogram ~ Quantiles

Bucket 1 Bucket 2 Bucket 3 Bucket 4 Bucket 5
 Count=9 Count=10 Count=10 Count=7 Count=9

CMU SCS

Today's Class

- History & Background
- Relational Algebra Equivalences
- Plan Cost Estimation
- Plan Enumeration

Faloutsos/Pavlo CMU SCS 15-415/615

40

CMU SCS

Query Optimization

- Bring query in internal form into “canonical form” (syntactic q-opt)
- Generate alternative plans.
 - ➔ – Single relation.
 - Multiple relations.
 - Nested sub-queries.
- Estimate cost for each plan.
- Pick the best one.

Faloutsos/Pavlo CMU SCS 15-415/615 41

CMU SCS

Plan Generation

SELECT *
FROM SAILORS
WHERE rating = 10

- What are our plan options?

Faloutsos/Pavlo CMU SCS 15-415/615 42

CMU SCS

Plan Generation

SELECT *
FROM SAILORS
WHERE rating = 10

- Sequential Scan
- Binary Search
 - *if sorted & consecutive*
- Index Search
 - *if an index exists*

Faloutsos/Pavlo CMU SCS 15-415/615 43

CMU SCS

Sequential Scan

SELECT *
FROM SAILORS
WHERE rating = 10

- B_R (worst case)
- $B_R / 2$ (on average, if we search for primary key)

Faloutsos/Pavlo CMU SCS 15-415/615 44

CMU SCS

Binary Search

SELECT *
FROM SAILORS
WHERE rating = 10

- $\sim \log(B_R) + SC(A,R)/F_R$
- Extra blocks are ones that contain qualifying tuples

#1
#2
#3
...
#B_R

Faloutsos/Pavlo CMU SCS 15-415/615 45

CMU SCS

Binary Search

SELECT *
FROM SAILORS
WHERE rating = 10

- $\sim \log(B_R) + SC(A,R)/F_R$
- Extra blocks are ones that contain qualifying tuples

We showed that estimating this is non-trivial.

#1
#2
#3
...
#B_R

Faloutsos/Pavlo CMU SCS 15-415/615 46

CMU SCS

Index Search

SELECT *
FROM SAILORS
WHERE rating = 10

- Index Search:
 - levels of index +
 - blocks w/ qual. tuples

Case#1: Primary Key
Case#2: Secondary key – clustering index
Case#3: Secondary key – non-clust. index

#1
#2
#3
...
#B_R

Faloutsos/Pavlo CMU SCS 15-415/615 47

CMU SCS

Index Search: Case #1

SELECT *
FROM SAILORS
WHERE rating = 10

- Primary Key
 - cost: $HT_i + 1$

#1
#2
#3
...
#B_R

Faloutsos/Pavlo CMU SCS 15-415/615 48

CMU SCS

Index Search: Case #2

SELECT *
FROM SAILORS
WHERE rating = 10

- Secondary key with clustering index:
 - cost: $HT_i + SC(A,R)/F_R$

Faloutsos/Pavlo CMU SCS 15-415/615 49

CMU SCS

Index Search: Case #3

SELECT *
FROM SAILORS
WHERE rating = 10

- Secondary key with non-clustering index:
 - cost: $HT_i + SC(A,R)$

Faloutsos/Pavlo CMU SCS 15-415/615 50

CMU SCS

Query Optimization

- Bring query in internal form into “canonical form” (syntactic q-opt)
- Generate alternative plans.
 - Single relation.
 - ➔ – Multiple relations.
 - Nested sub-queries.
- Estimate cost for each plan.
- Pick the best one.

Faloutsos/Pavlo CMU SCS 15-415/615 51

CMU SCS

Queries over Multiple Relations

- As number of joins increases, number of alternative plans grows rapidly
 - We need to restrict search space.
- Fundamental decision in System R:** only **left-deep join trees** are considered.

Faloutsos/Pavlo CMU SCS 15-415/615 52

CMU SCS

Queries over Multiple Relations

- **Fundamental decision in System R:** only left-deep join trees are considered.

Faloutsos/Pavlo CMU SCS 15-415/615 53

CMU SCS

Queries over Multiple Relations

- **Fundamental decision in System R:** only left-deep join trees are considered.
 - Allows for fully pipelined plans where intermediate results not written to temp files.
 - Not all left-deep trees are fully pipelined.

Faloutsos/Pavlo CMU SCS 15-415/615 54

CMU SCS

Queries over Multiple Relations

- **Enumerate the orderings**
 - Example: Left-deep tree #1, Left-deep tree #2...
- **Enumerate the plans for each operator**
 - Example: Hash, Sort-Merge, Nested Loop...
- **Enumerate the access paths for each table**
 - Example: Index #1, Index #2, Seq Scan...
- Use **dynamic programming** to reduce the number of cost estimations.

Faloutsos/Pavlo CMU SCS 15-415/615 55

CMU SCS

Dynamic Programming Example

Compute the cheapest flight PIT -> PVG
Solution: Compute partial optimal, left-to-right

Faloutsos/Pavlo CMU SCS 15-415/615 56

Q-Opt + Dynamic Programming

- Example: $R \bowtie S \bowtie T$

Faloutsos/Pavlo CMU SCS 15-415/615 57

Q-Opt + Dynamic Programming

- How to plan a query where **R** is sorted on **R.a**?
- Consider the following query:

```
SELECT *
FROM R, S, T
WHERE R.a = S.a AND S.b = T.b
ORDER BY R.a
```

Faloutsos/Pavlo CMU SCS 15-415/615 58

Q-Opt + Dynamic Programming

- $R \bowtie S \bowtie T$ order by R.a

Faloutsos/Pavlo CMU SCS 15-415/615 59

Candidate Plan Example

```
SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

- Enumerate relation orderings
- Enumerate join algorithm choices
- Enumerate access method choices

Faloutsos/Pavlo CMU SCS 15-415/615 60

CMU SCS

Candidate Plans

```
SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

1. Enumerate relation orderings:

Faloutsos/Pavlo CMU SCS 15-415/615 61

CMU SCS

Candidate Plans

```
SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

1. Enumerate relation orderings:

Faloutsos/Pavlo CMU SCS 15-415/615 62

CMU SCS

Candidate Plans

```
SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

1. Enumerate relation orderings:

Faloutsos/Pavlo CMU SCS 15-415/615 63

CMU SCS

Candidate Plans

```
SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

2. Enumerate join algorithm choices:

Faloutsos/Pavlo CMU SCS 15-415/615 64

Candidate Plans

SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid **AND** R.bid = B.bid

3. Enumerate access method choices:

Do this for the other plans.

Faloutsos/Pavlo CMU SCS 15-415/615 65

Candidate Plans

SELECT sname, bname, day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid **AND** R.bid = B.bid

4. Now we can estimate the cost of each plan.

Faloutsos/Pavlo CMU SCS 15-415/615 66

Query Optimization

- Bring query in internal form into “canonical form” (syntactic q-opt)
- Generate alternative plans.
 - Single relation.
 - Multiple relations.
 - ➔ – Nested sub-queries.
- Estimate cost for each plan.
- Pick the best one.

Faloutsos/Pavlo CMU SCS 15-415/615 67

Nested Sub-Queries

- The DBMS treats nested sub-queries in the where clause as functions that take parameters and return a single value or set of values.
- Two Approaches:
 - Rewrite to de-correlate and/or flatten them
 - Decompose nested query and store result to temporary table

Faloutsos/Pavlo CMU SCS 15-415/615 68

CMU SCS

Nested Sub-Queries: Rewrite

```

SELECT name FROM Sailors AS S
WHERE EXISTS (
  SELECT * FROM Reserves AS R
  WHERE S.sid = R.sid
  AND R.day = "2015-10-18"
)
    
```

↓

```

SELECT name
FROM Sailors AS S, Reserves AS R
WHERE S.sid = R.sid
AND R.day = "2015-10-18"
    
```

Faloutsos/Pavlo CMU SCS 15-415/615 69

CMU SCS

Nested Sub-Queries: Decompose

```

SELECT S.sid, MIN(R.day)
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid
AND R.bid = B.bid
AND B.color = 'red'
AND S.rating = (SELECT MAX(S2.rating)
                FROM Sailors S2)
GROUP BY S.sid
HAVING COUNT(*) > 1
    
```

For each sailor with the highest rating (over all sailors) and at least two reservations for red boats, find the sailor id and the earliest date on which the sailor has a reservation for a red boat.

Faloutsos/Pavlo CMU SCS 15-415/615 70

CMU SCS

Decomposing Queries

- For harder queries, the optimizer breaks up queries into blocks and then concentrates on one block at a time.
- Sub-queries are written to a temporary table that are discarded after the query finishes.

Faloutsos/Pavlo CMU SCS 15-415/615 71

CMU SCS

Decomposing Queries

```

SELECT S.sid, MIN(R.day)
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid
AND R.bid = B.bid
AND B.color = 'red'
AND S.rating = (SELECT MAX(S2.rating)
                FROM Sailors S2)
GROUP BY S.sid
HAVING COUNT(*) > 1
    
```

↑ Outer Block ↑ Nested Block

Faloutsos/Pavlo CMU SCS 15-415/615 72

CMU SCS

What Optimizers are Still Bad At

- Cardinality estimations are still hard.
- Problem Areas:
 - Prepared Statements
 - Correlated Columns
 - Plan Stability

More Information: Guy Lohman, *Is Query Optimization a "Solved" Problem?* SIGMOD Blog (April 2014), <http://svp.sigmod.org/?p=1075>

Faloutsos/Pavlo CMU SCS 15-415/615 73

CMU SCS

Prepared Statements

```

PREPARE myQuery (varchar, int) AS
SELECT cname, amt
FROM customer AS C, account AS A
WHERE C.acctno = A.acctno
AND C.state = $1
AND A.amt > $2
    
```

```

EXECUTE myQuery ('PA', 1000);
    
```

Faloutsos/Pavlo CMU SCS 15-415/615 74

CMU SCS

Prepared Statements

Prepared Statement Query Plan

```

    graph TD
      C[CUSTOMER] -- "state=$1 σ" --> J((⋈ acctno=acctno))
      A[ACCOUNT] -- "amt>$2 σ" --> J
      J -- "π cname, amt" --> P[Projection]
    
```

- What should be the join order for **CUSTOMER** and **ACCOUNT**?


Faloutsos/Pavlo CMU SCS 15-415/615 75

CMU SCS

Prepared Statements

- **Solution #1** – Rerun optimizer each time the query is invoked.
- **Solution #2** – Generate multiple plans for different values of the parameters.
- **Solution #3** – Choose the average value for a parameter and use that for all invocations.


Faloutsos/Pavlo CMU SCS 15-415/615 76



Correlated Columns

- We showed how selectivities are modeled as probabilities on whether a predicate on any given row will be satisfied.
- We then multiply these individual selectivities together.


Faloutsos/Pavlo CMU SCS 15-415/615 77



Correlated Columns

- Consider a database of automobiles:
 - # of Makes = 10, # of Models = 100
- And the following query:
 - **make="Honda" AND model="Accord"**
- With the independence and uniformity assumption, the selectivity is:
 - $1/10 * 1/100 = 0.001$
- But since only Honda makes Accords the real selectivity is $1/100 = 0.01$.


Faloutsos/Pavlo CMU SCS 15-415/615 78



Column Group Statistics

- Tell the DBMS that it should keep track of statistics for groups of columns together rather than just treating them all as independent variables.
- Only supported in commercial systems.


Faloutsos/Pavlo CMU SCS 15-415/615 79



Plan Stability

- We want to deploy a new version of a DBMS but need to make sure that there are no performance regressions.
- What if 99% of the query plans are faster on the newer DBMS version, but 1% are slower?


Faloutsos/Pavlo CMU SCS 15-415/615 80



Plan Stability

- **Solution #1** – Allow tuning hints in plans.
- **Solution #2** – Set the optimizer version number and migrate queries one-by-one to the new optimizer.
- **Solution #3** – Save query plan from old version and provide it to the new DBMS.


Faloutsos/Pavlo CMU SCS 15-415/615 81



Query Optimizer Overview

- System R:
 - Break query in query blocks
 - Simple queries (ie., no joins): look at stats
 - n-way joins: left-deep join trees; ie., only one intermediate result at a time
 - *Pros: smaller search space; pipelining*
 - *Cons: may miss optimal*
 - 2-way joins: NL and sort-merge


Faloutsos/Pavlo CMU SCS 15-415/615 82



Conclusions

- Ideas to remember:
 - Syntactic q-opt – do selections early
 - Selectivity estimations (uniformity, indep.; histograms; join selectivity)
 - Hash join (nested loops; sort-merge)
 - Left-deep joins
 - Dynamic programming

Faloutsos/Pavlo CMU SCS 15-415/615 83



Next Class

- How to refine database schemas through normalization and functional dependencies to remove redundancies and prevent loss data.

Faloutsos/Pavlo CMU SCS 15-415/615 84

CMU SCS

Equivalence of Expressions

- Q: How to prove a transformation rule?
- Use relational tuple calculus to show that LHS = RHS:

$$\sigma_p(R1 \cup R2) = \sigma_p(R1) \cup \sigma_p(R2)$$

LHS **RHS**

Faloutsos/Pavlo CMU SCS 15-415/615 85

CMU SCS

Equivalence of Expressions

$$\sigma_p(R1 \cup R2) = \sigma_p(R1) \cup \sigma_p(R2)$$

↓

$$t \in LHS \Leftrightarrow$$

$$t \in (R1 \cup R2) \wedge P(t) \Leftrightarrow$$

$$(t \in R1 \vee t \in R2) \wedge P(t) \Leftrightarrow$$

$$(t \in R1 \wedge P(t)) \vee (t \in R2) \wedge P(t) \Leftrightarrow$$

Faloutsos/Pavlo CMU SCS 15-415/615 86

CMU SCS

Equivalence of Expressions

$$\sigma_p(R1 \cup R2) = \sigma_p(R1) \cup \sigma_p(R2)$$

...

↓

$$(t \in R1 \wedge P(t)) \vee (t \in R2) \wedge P(t) \Leftrightarrow$$

$$(t \in \sigma_p(R1)) \vee (t \in \sigma_p(R2)) \Leftrightarrow$$

$$t \in \sigma_p(R1) \cup \sigma_p(R2) \Leftrightarrow$$

$$t \in RHS$$

QED

Faloutsos/Pavlo CMU SCS 15-415/615 87

CMU SCS

Equivalence of Expressions

- Q: How to disprove a rule?

$$\pi_A(R1 - R2) \neq (\pi_A(R1) - \pi_A(R2))$$

↓ ↓

A	B
Trump	squirrels

≠

A	B
Trump	knifefights

∅

A	B
Trump	squirrels

A	B
Trump	knifefights

Faloutsos/Pavlo CMU SCS 15-415/615 88