

UNCERTAINTY-AWARE QUERY EXECUTION TIME PREDICTION

Wentao Wu^{1,2}, Xi Wu², Hakan Hacigumus³, Jeff Naughton²

¹Microsoft Research

²University of Wisconsin-Madison

³Google

Query Execution Time Estimation

2

- Problem Definition
 - ▣ Given a query, estimate its running time *before* it runs.
 - ▣ Focus on OLAP style, *long-running* queries.

- Applications
 - ▣ Traditionally, cost-based query optimization.
 - ▣ Recently, database as a service (DaaS): admission control, query scheduling, system sizing, ...

Previous Work

3

- Single-Query Workload
 - ▣ [Ganapathi ICDE'09], [Xiong SoCC'11], [Akdere ICDE'12], [Li VLDB'12], [Wu ICDE'13]

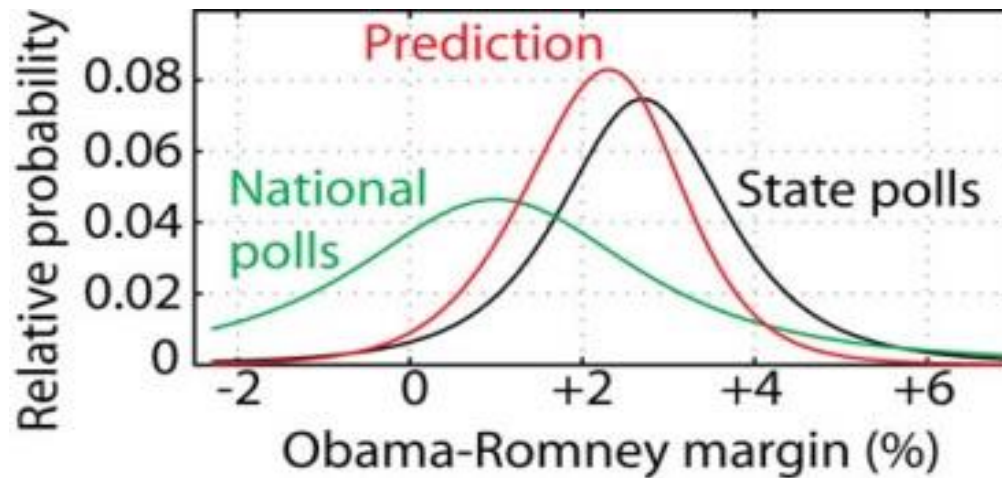
- Multi-Query Workload
 - ▣ [Ahmad EDBT'11], [Duggan SIGMOD'11], [Wu VLDB'13]

None of them is perfect, but none of them tried to quantify the *uncertainty* in the estimated query execution time.

Motivation

4

- Estimates are more useful with confidence intervals.



Measure Uncertainty: point estimate \Rightarrow distribution

Caveat

5

- What do we mean by “distribution of likely query execution times”?

Interpretation 1: If we run the query 100 times, what will be the distribution of its running times?



Interpretation 2: If we run the query now, what is the likelihood that it can finish between 100s and 200s?



Applications

6

- Query optimization
 - ▣ Least-Expected-Cost query optimization [Chu PODS'99]
 - ▣ Robust Query Optimization [Babcock SIGMOD'05]

- Query progress monitoring
 - ▣ Provide error bars for the “remaining” query running time.

- Database as a service
 - ▣ Distribution-based query scheduling [Chi VLDB'13].

Our Idea

7

□ PostgreSQL's cost model

$$C = n_s c_s + n_r c_r + n_t c_t + n_i c_i + n_o c_o$$

□ c 's: cost units

□ n 's: functions of *cardinality* estimates

| Cost Unit | Value |
|------------------------------|--------|
| c_s : seq_page_cost | 1.0 |
| c_r : rand_page_cost | 4.0 |
| c_t : cpu_tuple_cost | 0.01 |
| c_i : cpu_index_tuple_cost | 0.005 |
| c_o : cpu_operator_cost | 0.0025 |

In our previous work [Wu ICDE'13], we proposed a framework that *calibrates* the c 's and *refines* the n 's to get *better* query execution time estimates.


Our Idea (Cont.)

8



$$t = n_s c_s + n_r c_r + n_t c_t + n_i c_i + n_o c_o$$

View the c 's and the n 's as *random variables* rather than constants!



Pr (c)



Pr (n)



$$t = n \cdot c$$



Pr (t)

Outlines

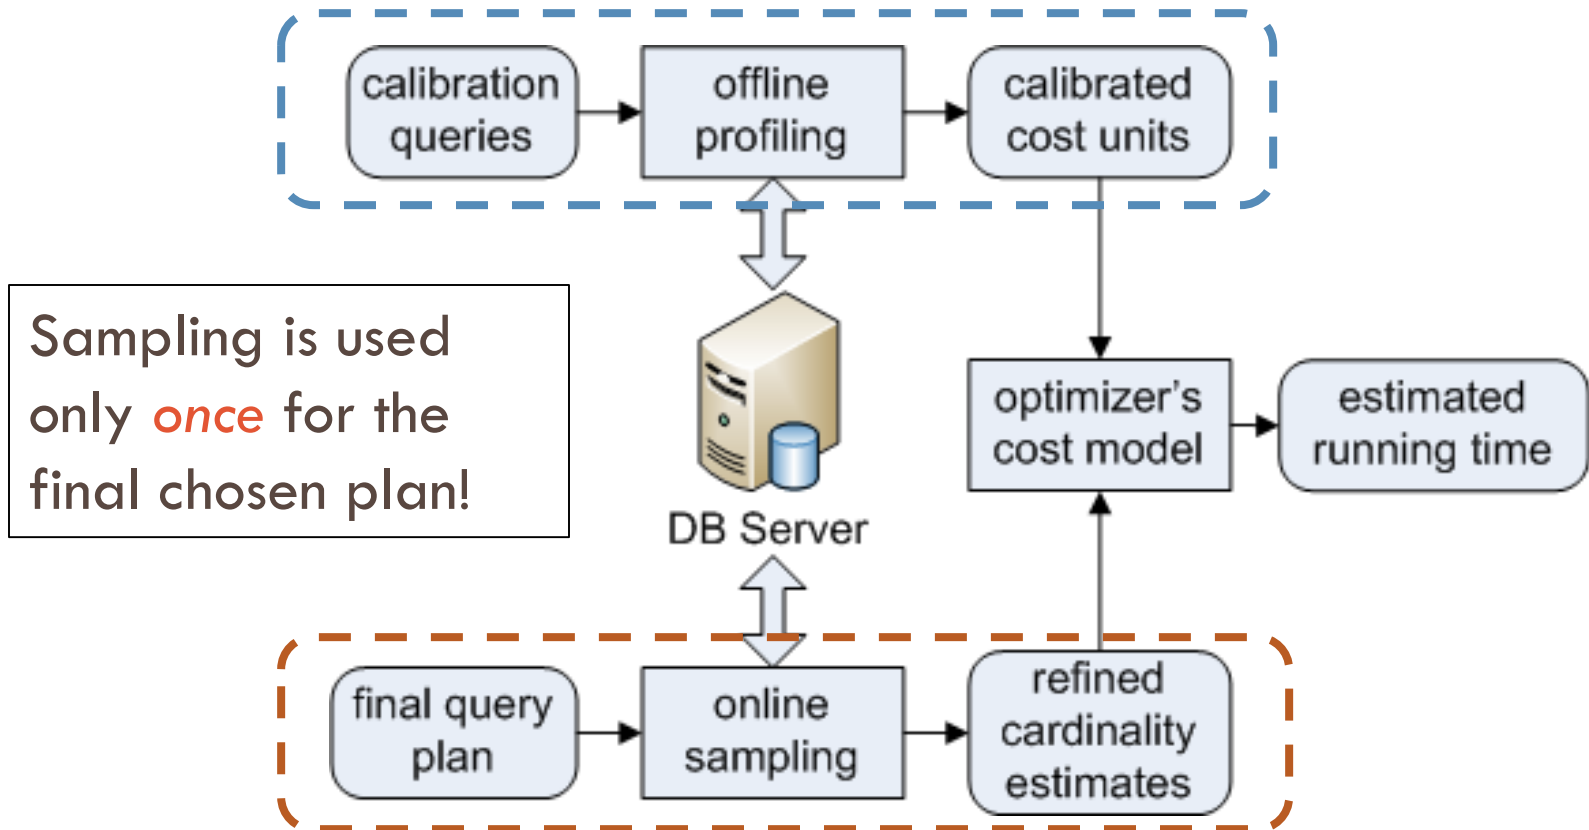
9

- The calibration framework
- Distributions of the c 's
- Distributions of the n 's
- Summary

The Calibration Framework

10

- Calibrate the c 's: use calibration queries.



- Refine the n 's: refine cardinality estimates.

Outlines

11

- The calibration framework
- Distributions of the c 's
- Distributions of the n 's
- Summary

Calibrate The c's

12

- Basic idea (an example)
 - ▣ Want to know the true values of c_t and c_o via calibration queries.

```
q1: select * from R
q2: select count(*)
      from R
```

R in memory



| Cost Unit |
|------------------------------|
| c_s : seq_page_cost |
| c_r : rand_page_cost |
| c_t : cpu_tuple_cost |
| c_i : cpu_index_tuple_cost |
| c_o : cpu_operator_cost |

$$t_1 = c_t \cdot n_t$$
$$t_2 = c_t \cdot n_t + c_o \cdot n_o$$

- General case
 - ▣ k cost units (i.e., k **unknowns**) \Rightarrow k queries (i.e., k **equations**)
 - ▣ $k = 5$ in the case of PostgreSQL

Calibration Queries For PostgreSQL

13

Isolate the unknowns and solve them *one per equation!*

q_1 : select * from R

R in memory

$$t_1 = c_t \cdot n_{t1}$$

q_2 : select count(*) from R

R in memory

$$t_2 = c_t \cdot n_{t2} + c_o \cdot n_{o2}$$

q_3 : select * from R where R.A < a (R.A with an Index)

R in memory

$$t_3 = c_t \cdot n_{t3} + c_i \cdot n_{i3} + c_o \cdot n_{o3}$$

q_4 : select * from R

R on disk

$$t_4 = c_s \cdot n_{s4} + c_t \cdot n_{t4}$$

q_5 : select * from R where R.B < b (R.B *unclustered* Index)

R on disk

$$t_5 = c_s \cdot n_{s5} + c_r \cdot n_{r5} + c_t \cdot n_{t5} + c_i \cdot n_{i5} + c_o \cdot n_{o5}$$

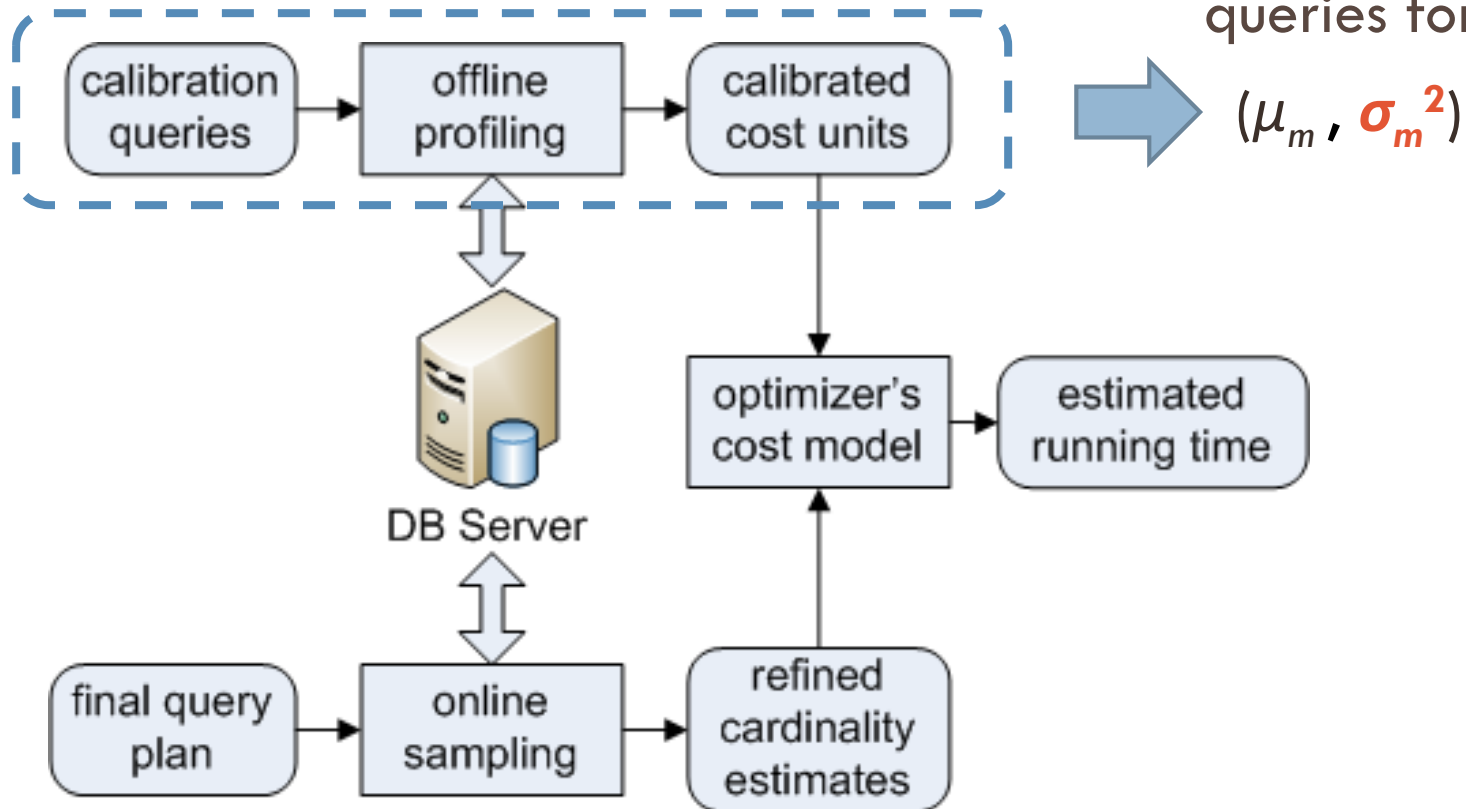
For each c , use *multiple* queries and take the *average*.

Distributions of the c's

14

- Assumption: $c \sim N(\mu, \sigma^2)$

m calibration queries for each c



Outlines

15

- The calibration framework
- Distributions of the c 's
- Distributions of the n 's
- Summary

Refine The n 's

16

- The n 's are *functions* of N 's (i.e., input cardinalities).

Example 1 (In-Memory Sort)

$$sc = [2 \cdot N_t \cdot \log N_t] \cdot c_o + tc \text{ of child}$$
$$rc = c_t \cdot N_t$$

Example 2 (Nested-Loop Join)

$$sc = sc \text{ of outer child} + sc \text{ of inner child}$$

$$rc = c_t \cdot N_t^o \cdot N_t^l + N_t^o \cdot rc \text{ of inner child}$$

sc : start-cost rc : run-cost tc : total-cost N_t : # of input tuples

Distributions of the n 's

17

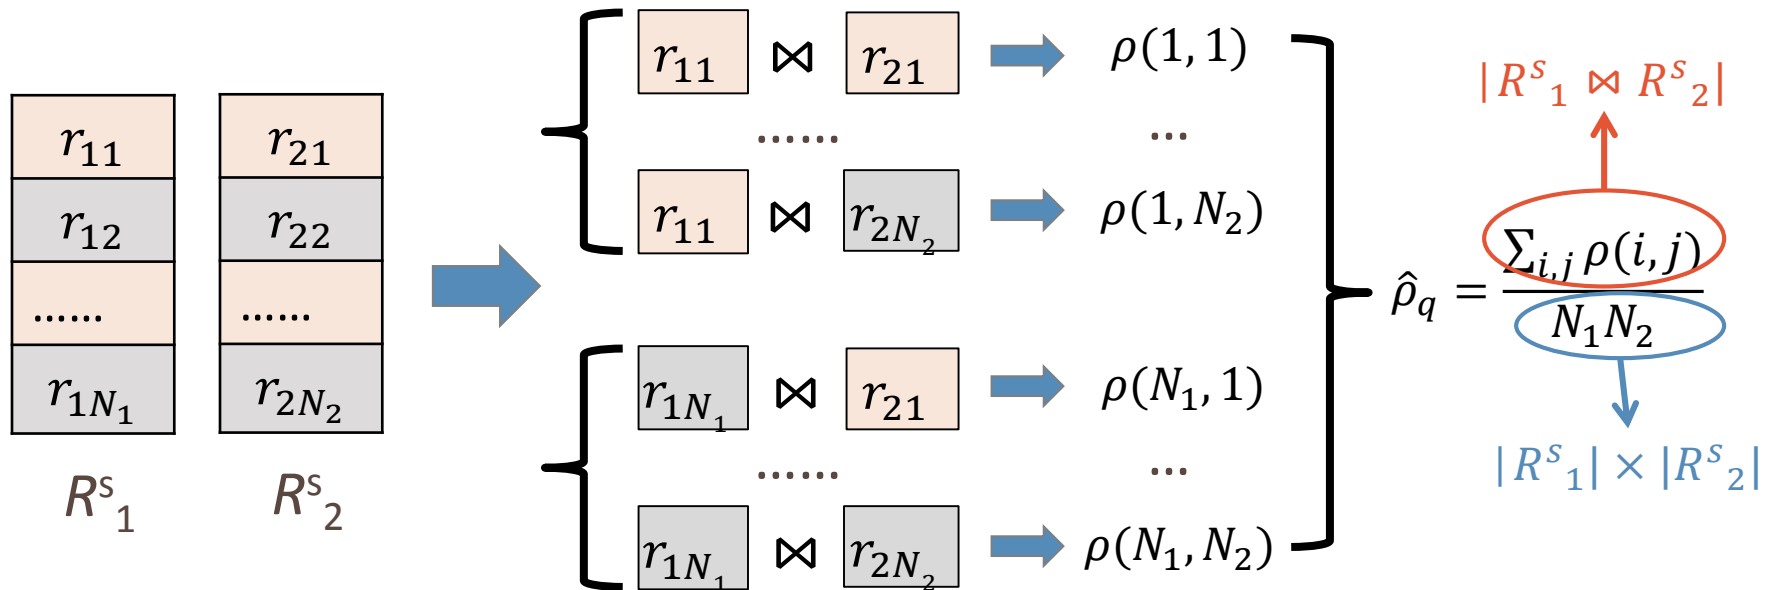
- We need to model two quantities:
 - ▣ The selectivities;
 - ▣ The cost functions for different physical operators.
- Using mathematics we get distributions of the n 's.

A Sampling-Based Selectivity Estimator

- Estimate the selectivity ρ_q of a join query $q = R_1 \bowtie R_2$.

[Haas et al., J. Comput. Syst. Sci. 1996]

Do a “cross product” over the samples: $\rho(i, j) = 0$ or 1 .

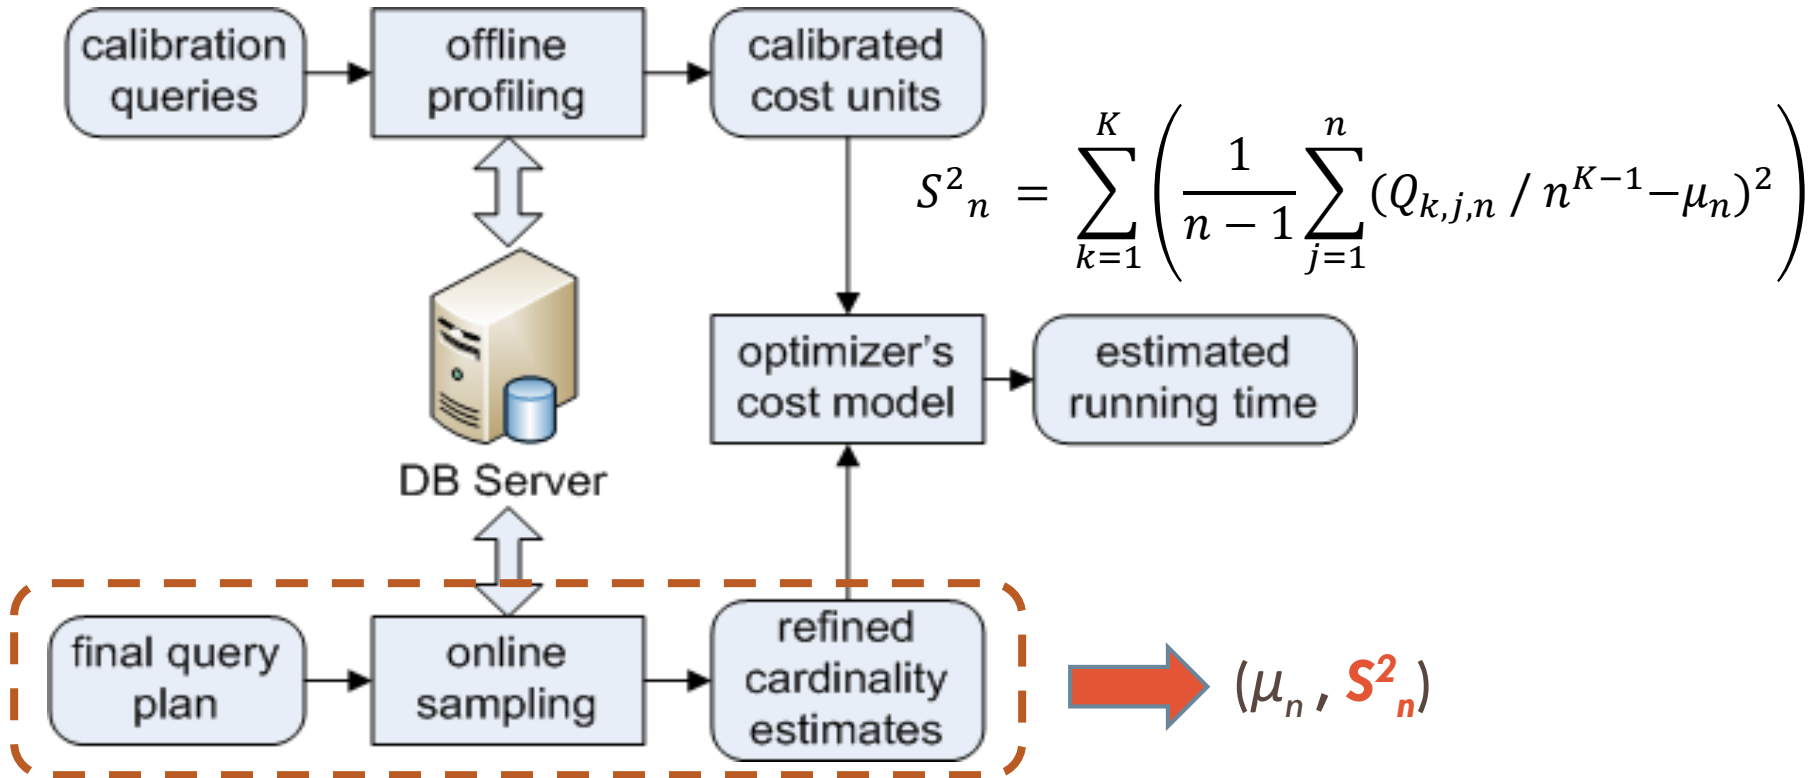


The ρ 's are not independent, but the estimator $\hat{\rho}_q$ is still *unbiased* and *strongly consistent*.

Distributions of Selectivities

19

- Selectivity $\sim N(\mu_n, S_n^2)$: by the Central Limit Theorem.



Modeling Cost Functions

20



(Different Implementations)

Nested-Loop Join

Generic Cost Function: $f = a_0 N_l N_r + a_1 N_l + a_2 N_r + a_3$

N_l and N_r are the left and right input cardinality of the operator.

Distributions of the n 's and t

21

- The n 's are *asymptotically* Gaussian.
 - More samples \Rightarrow more close to Gaussian
- The running time t is also *asymptotically* Gaussian.

Example: Nested Loop Join

$$t = n_r c_r + n_t c_t \quad \Rightarrow \quad t \sim N(E[t], \text{Var}[t])$$

$$\begin{cases} n_r = a_0 N_l N_r + a_1 N_l + a_2 N_r + a_3 \\ n_t = b_0 N_l N_r + b_1 N_l + b_2 N_r + b_3 \end{cases} \quad \Rightarrow \quad n_r \text{ and } n_t \text{ are *not* independent!}$$

Should consider *covariances* when computing $\text{Var}[t]$!

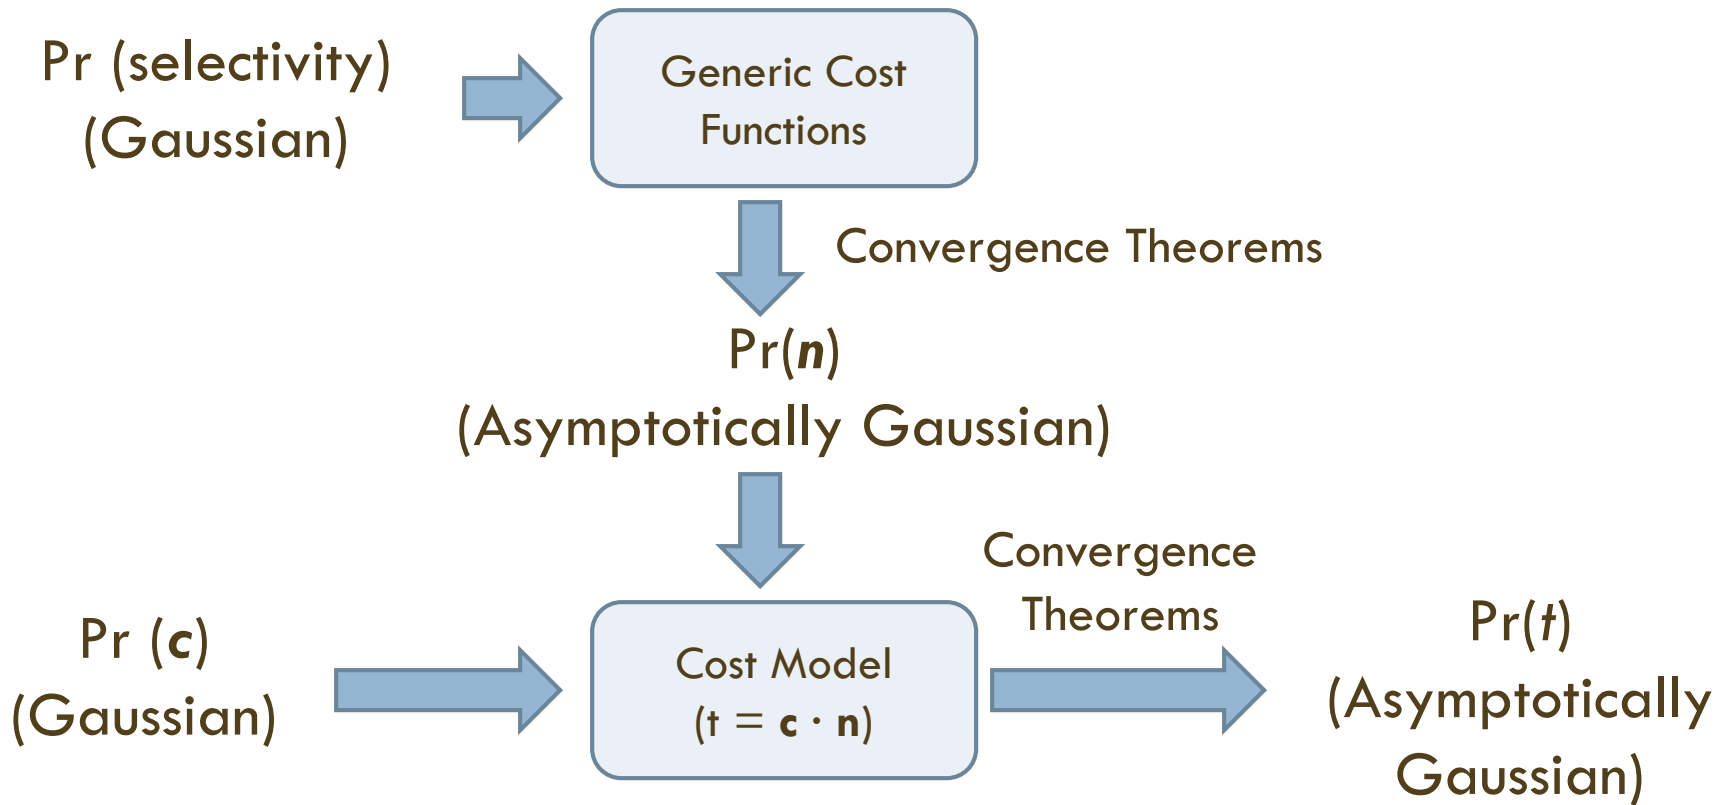
Outlines

22

- The calibration framework
- Distributions of the c 's
- Distributions of the n 's
- Summary

Put It Together (Review)

23




Experimental Evaluation

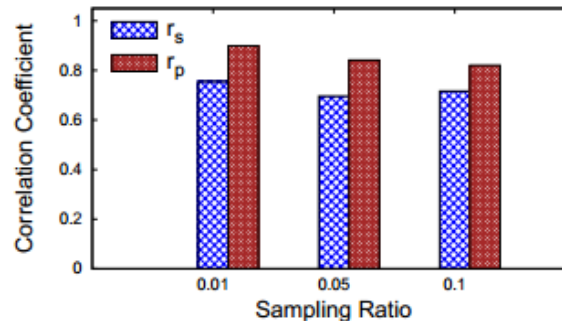
24

- The idea: larger variances \Rightarrow larger estimation errors


We thus measure the *correlation* between the two.



(a) MICRO, Uniform 1GB, PC2



(b) SELJOIN, Uniform 1GB, PC1



(c) TPCH, Skewed 10GB, PC1

We observed *strong* correlations (i.e., correlation coefficient > 0.7) on almost all the queries tested in our experiments.

Q & A

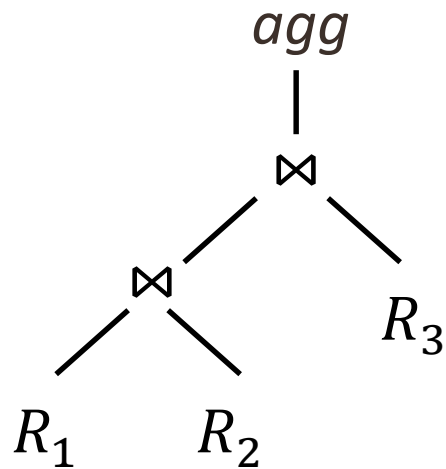
25

□ Thank you😊

The Cardinality Refinement Algorithm (Example)

26

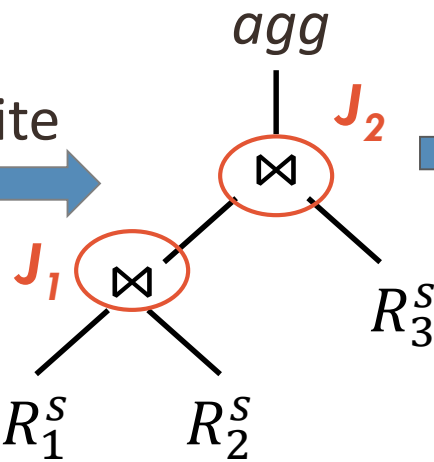
Plan of q :



$$J_1 = R_1 \bowtie R_2$$

$$J_2 = R_1 \bowtie R_2 \bowtie R_3$$

Rewrite



Run

$$\hat{\rho}_{J_1} = \frac{|R_1^S \bowtie R_2^S|}{|R_1^S| \times |R_2^S|}$$

$$\hat{\rho}_{J_2} = \frac{|R_1^S \bowtie R_2^S \bowtie R_3^S|}{|R_1^S| \times |R_2^S| \times |R_3^S|}$$

Reuse

R_1^S, R_2^S, R_3^S are samples (as tables) of R_1, R_2, R_3 .

For *agg*, use PostgreSQL's estimates based on the *refined* input estimates from J_2 .

Distributions of Selectivities (Cont.)

Implementation of S^2_n in PostgreSQL

$$S^2_n = \sum_{k=1}^K \left(\frac{1}{n-1} \sum_{j=1}^n (Q_{k,j,n} / n^{K-1} - \mu_n)^2 \right)$$

If: $r_{11} \bowtie r_{21} \in R^{s_1} \bowtie R^{s_2}$
 Then: $++Q_{1,1,n}$
 $++Q_{2,1,n}$

Example: $K = 2$ (i.e., $R_1 \bowtie R_2$)

$$\begin{cases} Q_{1,j,n} = |\{r_{1j}\} \bowtie R^{s_2}| \\ Q_{2,j,n} = |R^{s_1} \bowtie \{r_{2j}\}| \end{cases}$$

r_{1j} and r_{2j} are the j -th row of R^{s_1} and R^{s_2} .

