Sub-optimal Join Order Classification by L1-error

Yesdaulet Izenov, Asoke Datta, Brian Tsan and Florin Rusu
University of California, Merced
yizenov, adatta2, btsan, frusu@ucmerced.edu

Q-error Limitations

Goal: classification of sub-optimal join orders

Q-error is a poor indicator of sub-optimal join orders:
- too loose as a bound
- does not consider enumeration algorithms
- cardinality estimation errors matter equally

Q-error Bound (JOB queries)

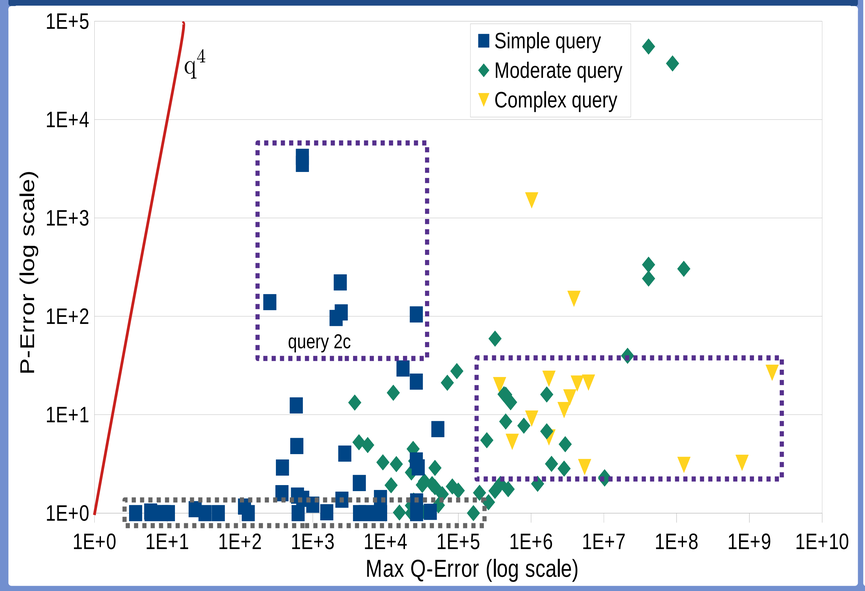

Query Plans (query 2c)

2-way joins (COST + CARDINALITY)

index ρ	1	2	3	4	5
sub-plan	$cn \bowtie mc$	$k \bowtie mk$	$mc \bowtie t$	$mk \bowtie t$	$mc \bowtie mk$
true Y	388	41.8K	2.6M	4.5M	34.9M
est. $\hat{Y}$	973	20	1.5M	2.7M	13.8M
Q-error q	2.51	2,092	1.70	1.70	2.53

3-way joins (COST + CARDINALITY)

index ρ	1	2	3	4	5
sub-plan	$cn \bowtie mc \bowtie t$	$cn \bowtie mc \bowtie mk$	$k \bowtie mk \bowtie t$	$k \bowtie mk \bowtie mc$	$mc \bowtie t \bowtie mk$
true Y	388+388	388+1,588	41.8K+41.8K	41.8K+148.6K	2.6M+34.9M
est. $\hat{Y}$	973+973	973+8,739	20+20	20+104	1.5M+2.7M
Q-error q	2.51	5.50	2,092	1,428.38	12.69

4-way joins (COST + CARDINALITY)

index ρ	1	2	3
sub-plan	$cn \bowtie mc \bowtie mk \bowtie k$	$cn \bowtie mc \bowtie t \bowtie mk$	$k \bowtie mk \bowtie t \bowtie mc$
true Y	1,976+4	776+1,588	83.7K+148.6K
est. $\hat{Y}$	124+1	1,946+8,739	40+104
Q-error q	4	5.50	1,428.38

5-way joins (COST + CARDINALITY)

index ρ	1
sub-plan	$cn \bowtie mc \bowtie mk \bowtie k \bowtie t$
true Y	1,980+4
est. $\hat{Y}$	125+1
Q-error q	4

Plan	Join Order	$\mathcal{P}\text{-}Cost(\mathcal{P}, Y)$	$\mathcal{P}\text{-}Cost(\mathcal{P}, \hat{Y})$
$\mathcal{P}_{opt}$	$cn \bowtie mc \bowtie mk \bowtie k \bowtie t$	1,980	9,713
$\mathcal{P}_{pg}$	$k \bowtie mk \bowtie mc \bowtie cn \bowtie t$	190.4K	125

L1-error

Sub-query Relative Order: estimation errors do not matter while sub-query relative order is preserved based on cardinality estimates
Sub-query Weights: the larger the weight, the greater the difference between two sub-queries
Small Cardinalities: sub-query with smaller cardinality is likely to be selected by the plan search algorithm (e.g. greedy algorithm)
Join Size Importance: early-stage multi-way joins are more critical and tend to be more accurate than joins at later stages

Dataset and Workload

IMDB Benchmark Dataset: real-world dataset containing correlations and non-uniform data distributions
Join Order Benchmark: challenging realistic workload — 113 queries
- 45 Simple queries with 4-9 join predicates
- 53 Moderate queries with 10-19 join predicates
- 15 Complex queries with 20-28 join predicates

Classification Results (JOB queries)

		Predicted		
		Sub-optimal (Positive)	Optimal (Negative)	
Actual	Sub-optimal (Positive)	94 (TP)	1 (FN)	95
	Optimal (Negative)	7 (FP)	11 (TN)	18
		101	12	

PostgreSQL estimates. Exhaustive Plan Search.

Simple Cost Function (query 2c)

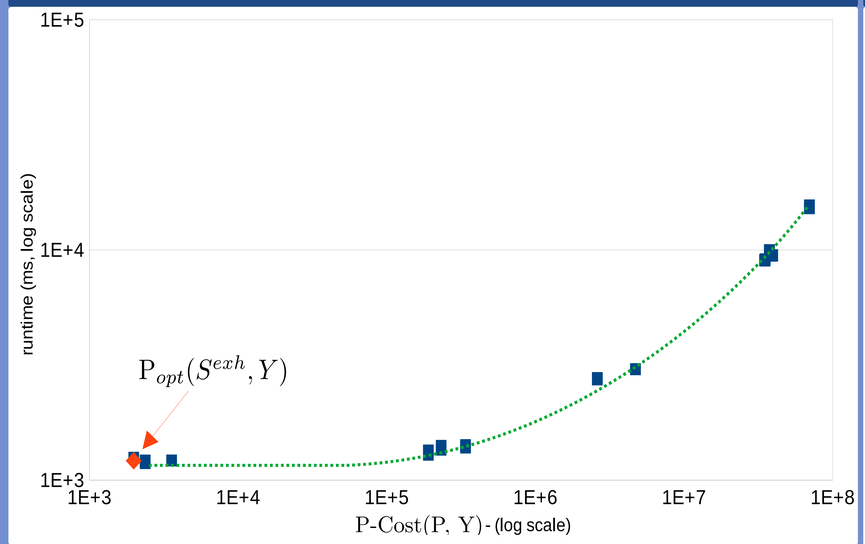

Cardinality Accuracy (JOB queries)

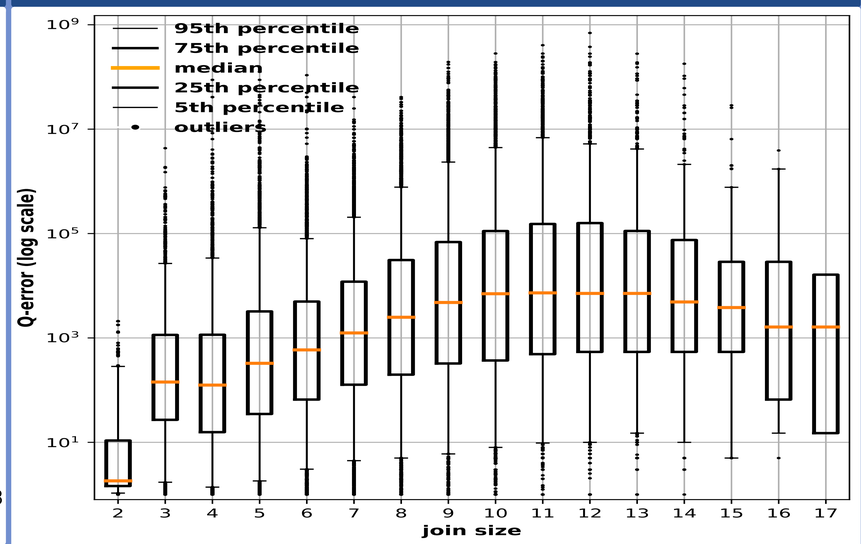

Join Size Importance (JOB queries)