Turbo-charging Picasso Diagram Generation

Database Systems Laboratory
Indian Institute of Science

Background

Query Template
+ Engine Choice
+ Diagram Resolution

Picasso [VLDB05]
→ Plan Diagram
→ Cost Diagram
→ Cardinality Diagram

DB Engine
DB2
Oracle
SQL Server
Sybase
PostgreSQL

Highly irregular plan boundaries

[TPCH – QT8]

Increases to 90 plans with 300x300 grid !

Resolution : 100 x100
# of plans : 76

Intricate Complex Patterns

- Generating 2D Optimizer Diagrams at resolution 1000, or 3D at resolution 100, requires 10^6 optimizations
- Cost of 1 optimization: ~ **0.5 sec**   Running time: ~ **1 WEEK !**

Can we obtain an accurate approximation in reasonable time?

Approximation Quality Metrics

- **Plan Identity Error (ε_I) :** % of plans which remained unidentified in approximation relative to the true plan diagram

$$\epsilon_I = \frac{|P| - |A|}{|P|} \times 100$$

- **Plan Location Error (ε_L) :** % of points assigned wrong plan in approximation relative to the true plan diagram

$$\epsilon_L = \frac{|p_A(q) \neq p_P(q)|}{m} \times 100$$

- User specifies bounds on these errors θ_I and θ_L

Query Optimizers

Class III
Class II
Class I
Optimal Plan
Plan Costing
Plan Rank List

Solution Framework

- Database → Random Sampling
+ Database → Non-Intrusive → Grid Processing
+ Database → Intrusive → Planfill with PRL

Grid Sampling

- Partition the Selectivity Space into smaller rectangles, optimize the **corners**.
- Process **middle** points of each edge
  - If end points are the same plan, assign this plan to the middle also
  - Else explicitly optimize the point
- Process **center** of each rectangle
- Iteratively partition until 1x1 box is reached

PQO principle: *If two points in a query parameter space have the same optimal plan, then this plan is optimal at all points on the straight line joining them .*

Attribute 2 / Attribute 1

Plan tree difference can be used as an indicator of "Plan Density"

Jaccard Distance between two trees

Let two plan trees T_i and T_j has $|T_i|$ and $|T_j|$ number of nodes respectively and $|T_i \cap T_j|$ denote the number of similar nodes between them. Then **Plan Density factor** is given by $\rho(T_i, T_j)$

$$\rho = 1 - \frac{T_i \cap T_j}{T_i \cup T_j}$$

More difference / Less difference

Matching nodes are colored white

Indication of lowest Plan Density

Indication of highest Plan Density

Planfill with 'Plan Rank List'

1. Optimize point q.
   - Let P_1 be the best and P_2 be the second best plan at q with costs Cost_P_1(q) and Cost_P_2(q), respectively   [Cost_P_1(q) < Cost_P_2(q)]

2. Foreign-plan-cost P_1 at points q' sweeping outward in first quadrant of q (i.e. increasing selectivity)
   - If Cost_P_1(q') ≤ Cost_P_2(q), assign P_1 to q'
   - Else repeat the process from Step 1 for q'

Planfill ensures that the diagram is generated with zero identity and location error.

Contributions

Optimizer Class	GenTime wrt Exhaustive [θ =10%]	GenTime wrt Exhaustive [Error θ =1%]
Class I (OP)	1% - 15%	15% - 40%
Class II (OP + FPC)	1% - 10%	15% - 40%
Class III (OP + FPC + PRL)	1% - 5%	1% -10%

- x10 improvement in generation time for typical case [θ =10%]
- In fact, can obtain guaranteed perfect [θ = 0%] diagrams in Class III with less than 10% overheads !
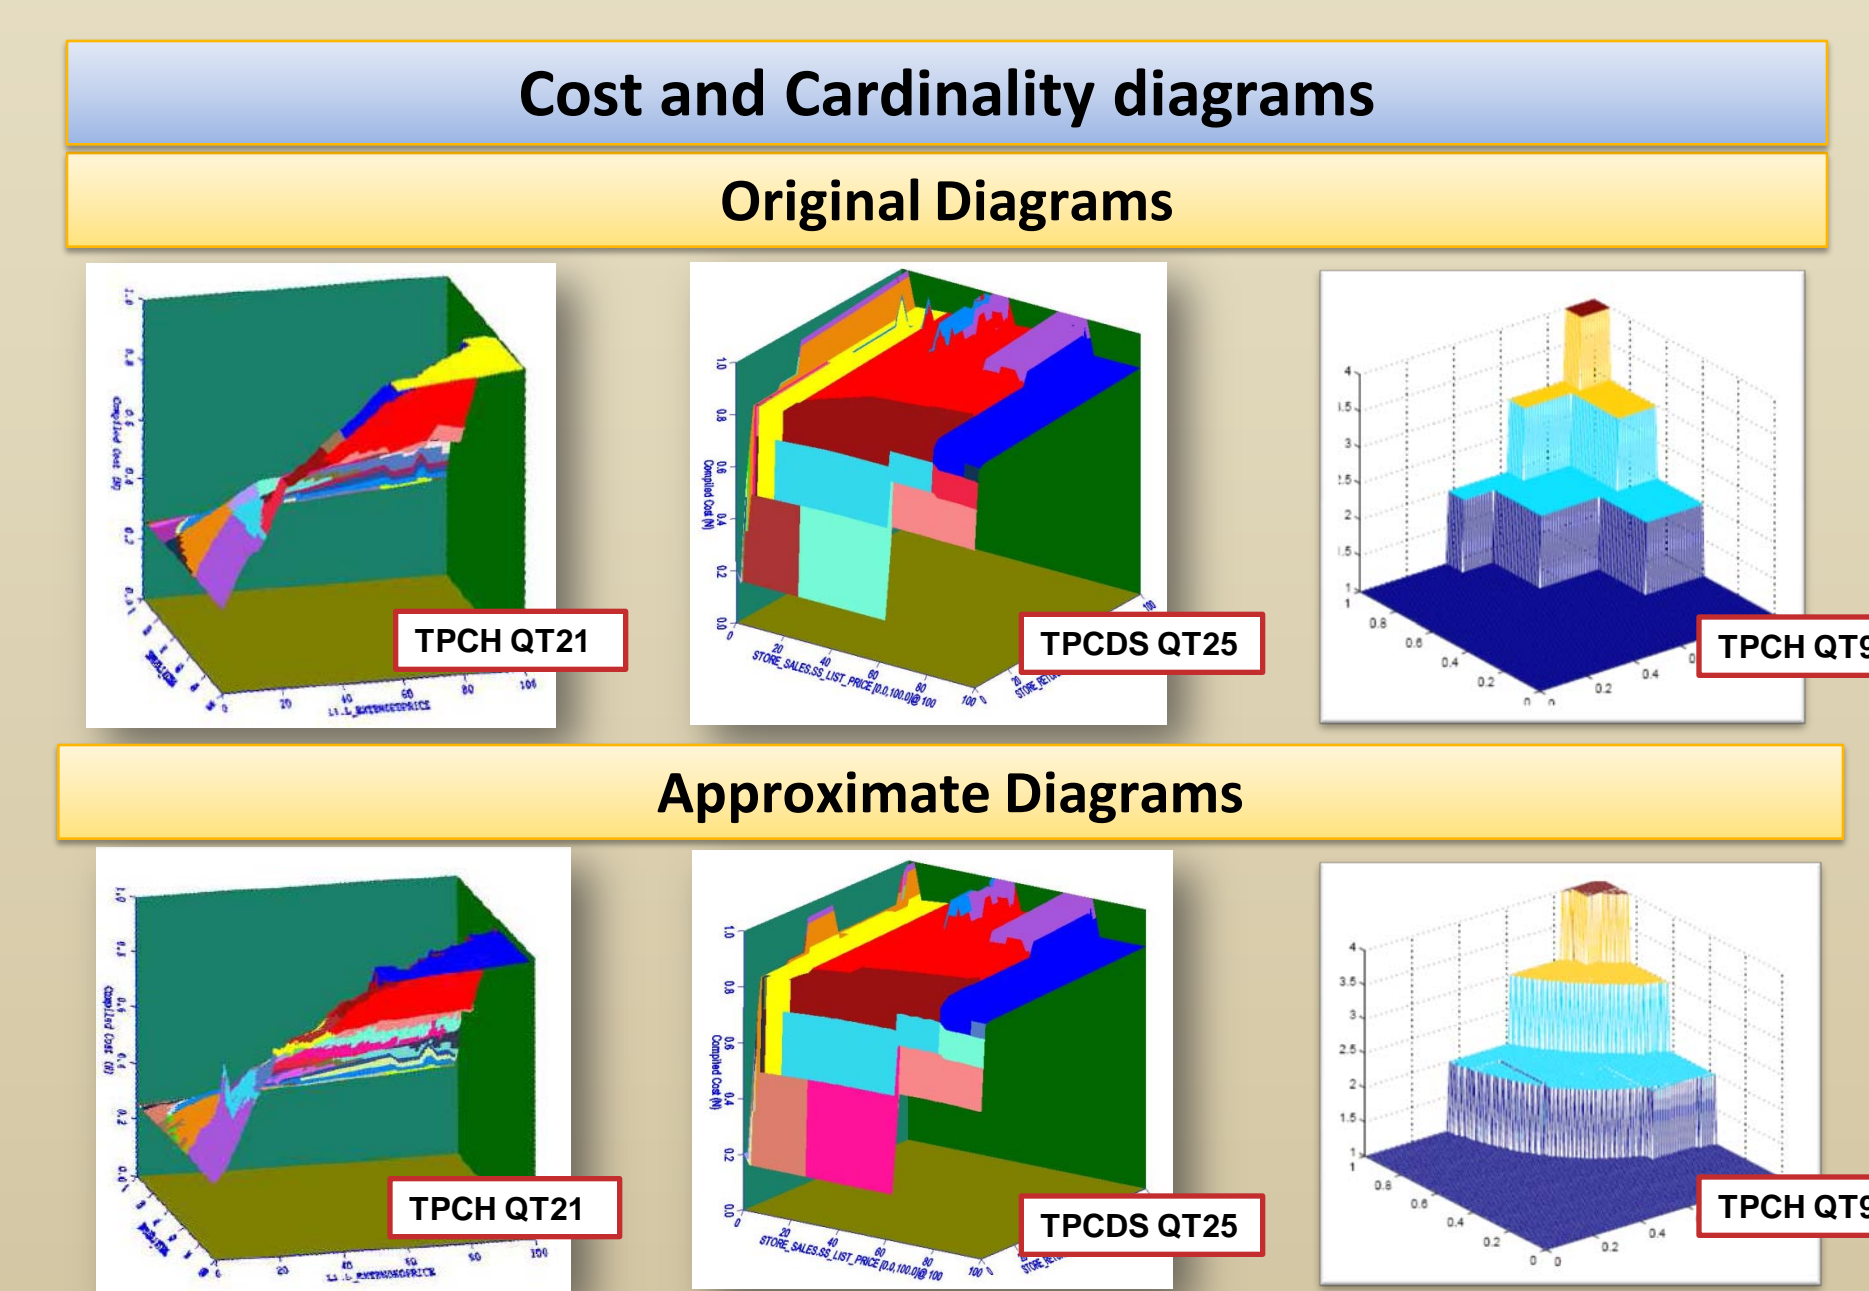- Cost and cardinality diagrams are also approximated using interpolation techniques.
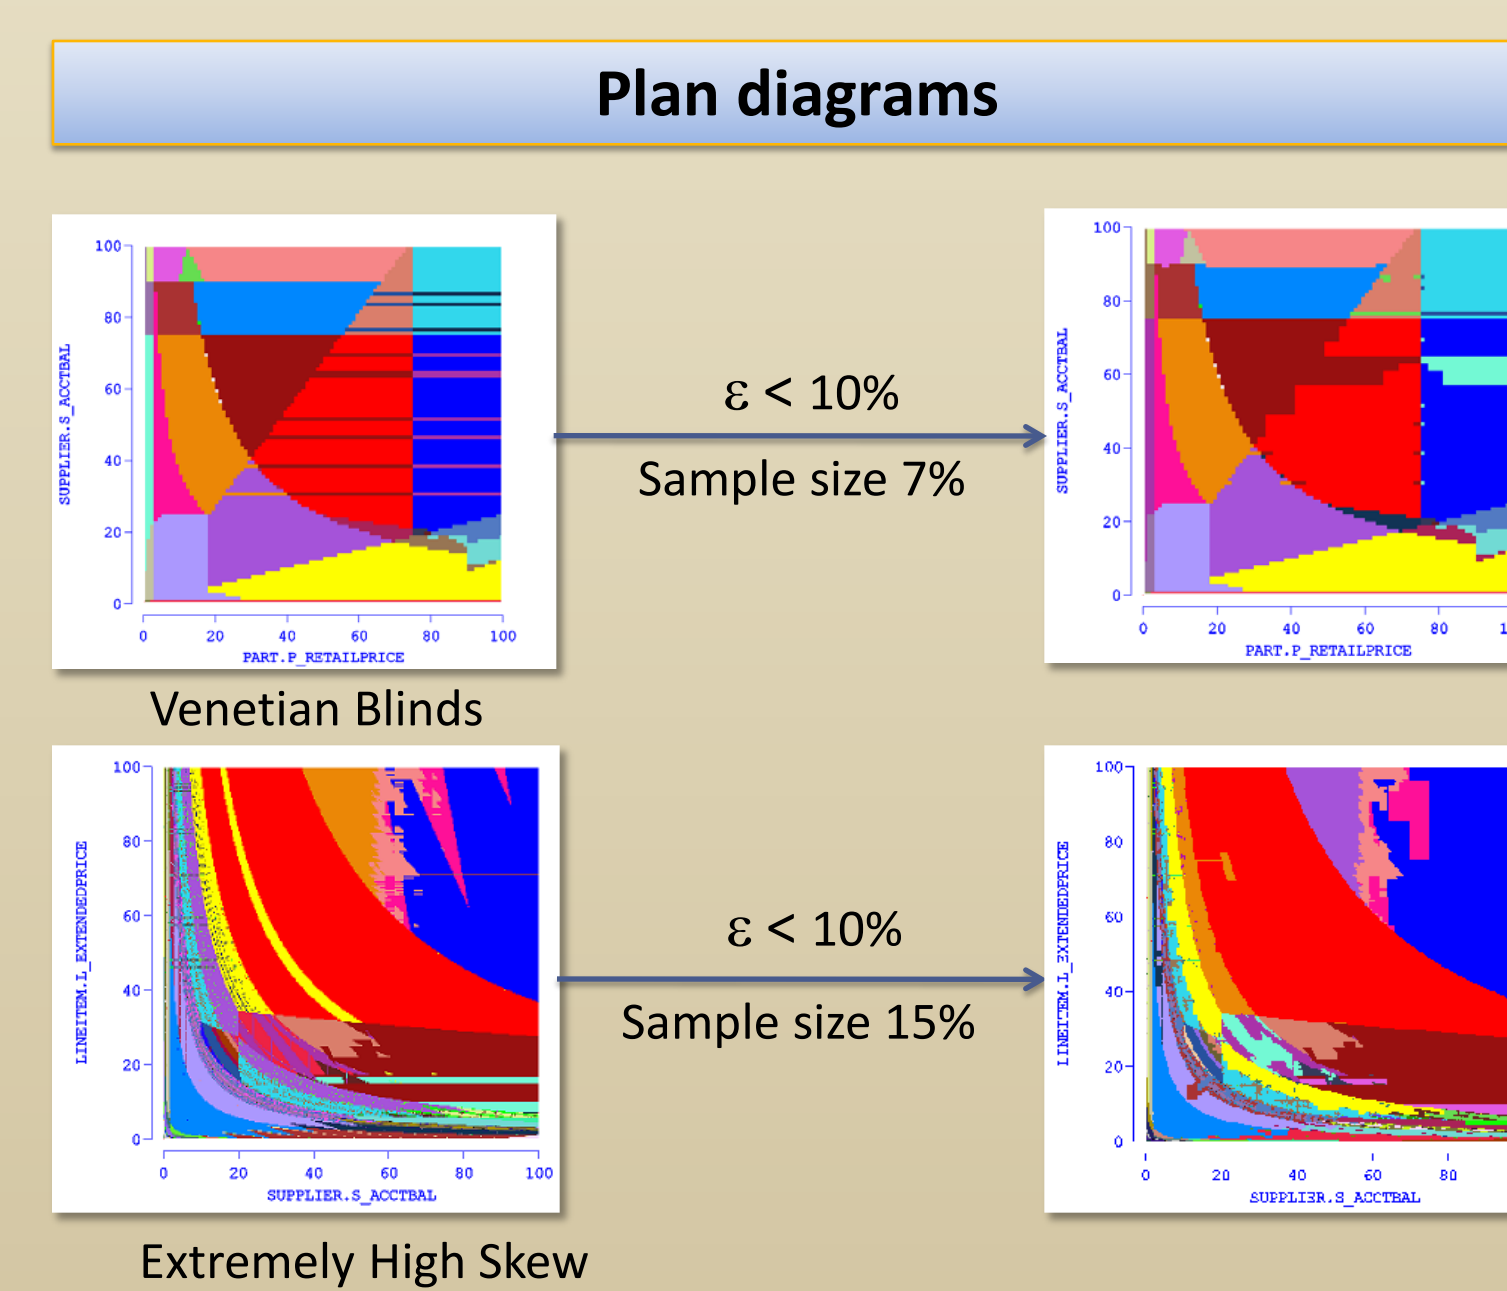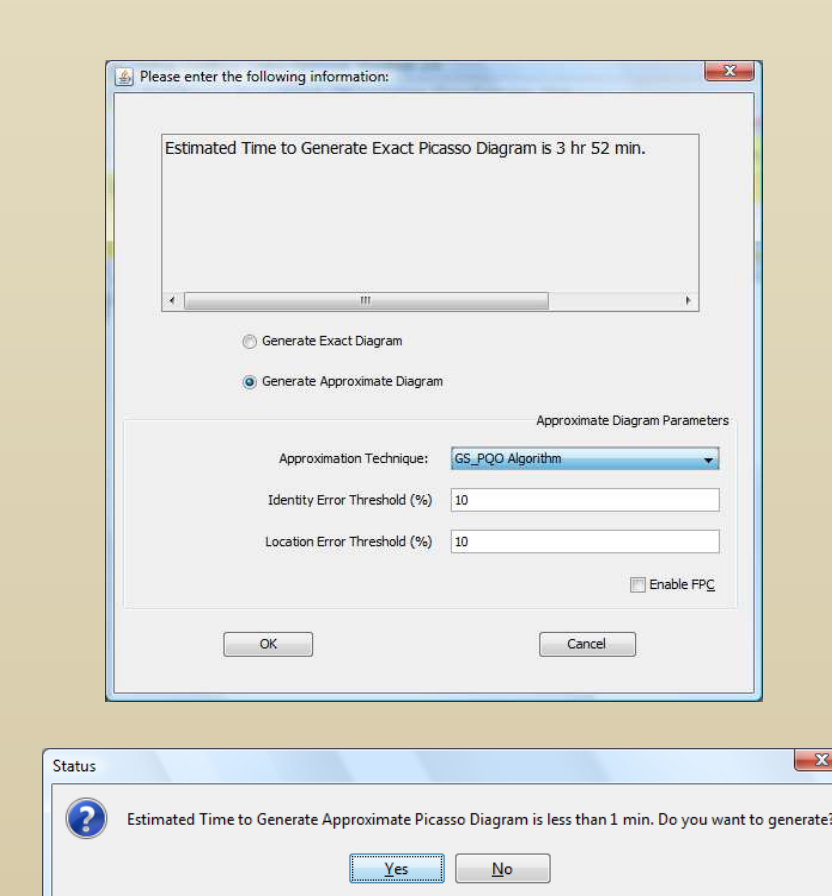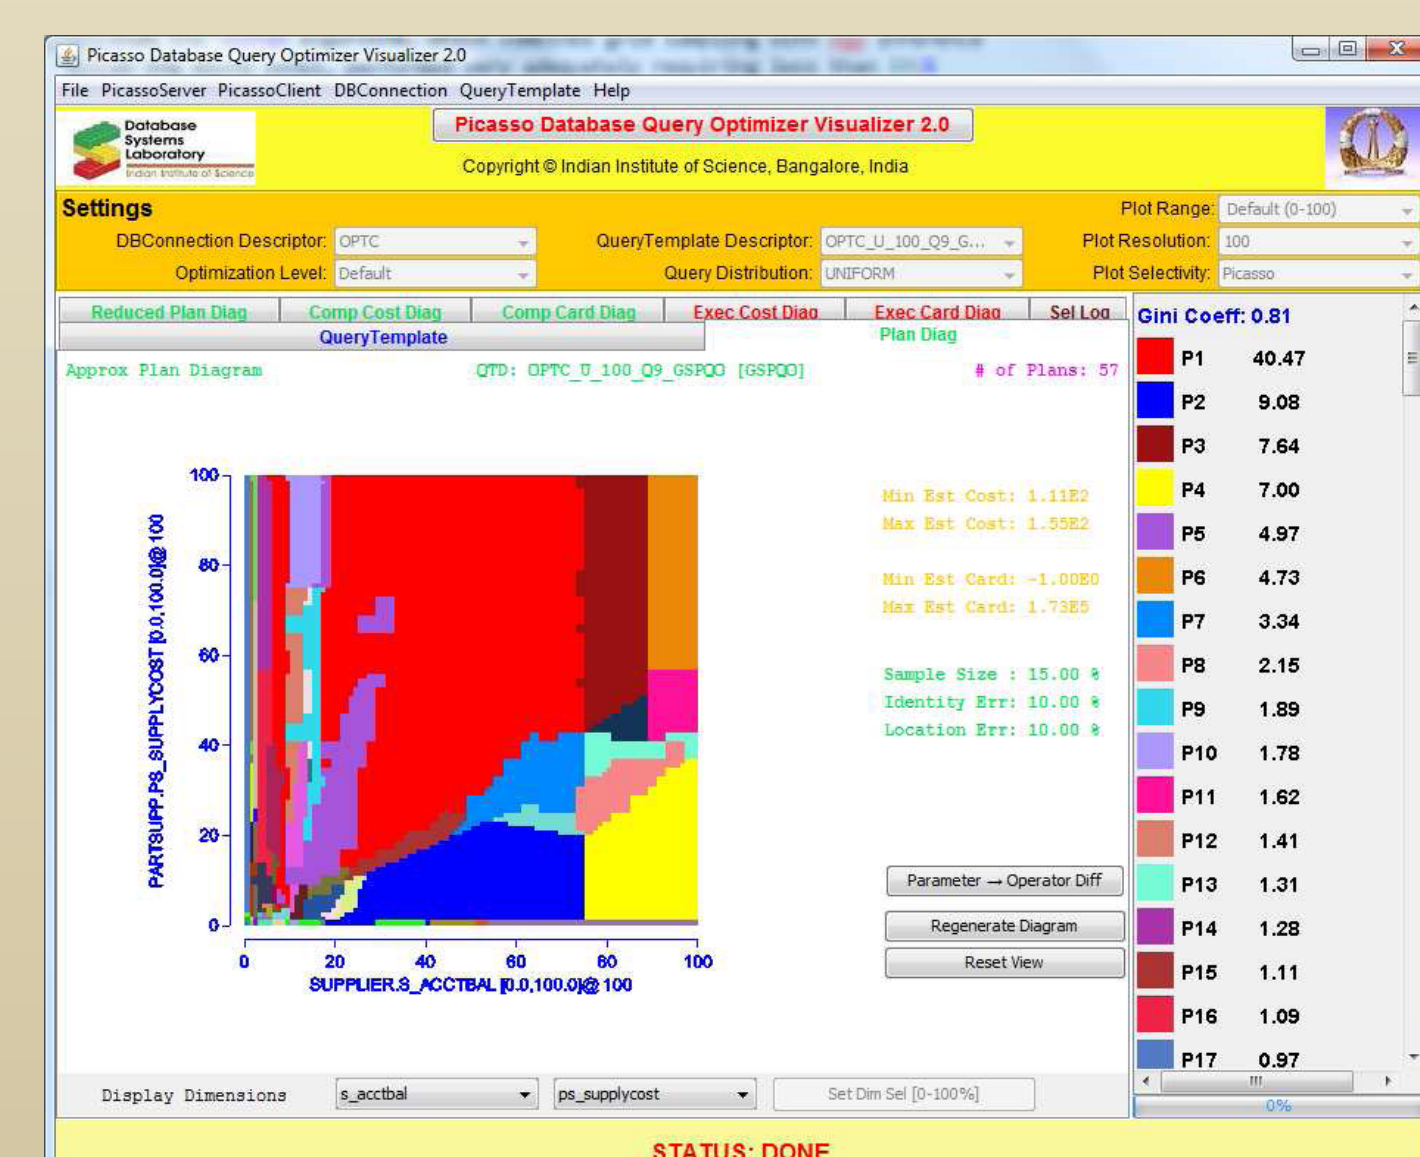
Implementation and results

- **Picasso 2.0**
  Class I and II optimizers were implemented.
- **FPC in PostgreSQL 8.3.6**
  Followed a depth-first search approach to re-cost a particular access-path on different selectivity.
- **Second Best Plan**
  Implemented the *Additive second best plan search* algorithm.

Plan diagrams

ε < 10%
Sample size 7%

Venetian Blinds

ε < 10%
Sample size 15%

Extremely High Skew

Cost and Cardinality diagrams

Original Diagrams

TPCH QT21   TPCDS QT25   TPCH QT9

Approximate Diagrams

TPCH QT21   TPCDS QT25   TPCH QT9

"Efficiently Approximating Query Optimizer Plan Diagrams", A. Dey, S. Bhaumik, Harish D. and J. Haritsa, Proc. of 34th Intl. Conf. on Very Large Data Bases (VLDB), Auckland, New Zealand, August 2008, pgs. 1325-1336.