Tutorial: The Strange and Wondrous Ways of Industrial-strength Database Query Optimizers

Jayant R. Haritsa

Database Systems Lab
Indian Institute of Science, Bangalore, India
haritsa@dsl.serc.iisc.ernet.in

Modern relational database systems incorporate a query optimizer module to identify the most efficient strategy, or plan, to execute the declarative SQL queries submitted by users. Optimization is a mandatory exercise since the difference between the cost of the best plan and a random choice could be in orders of magnitude. The role of query optimizers has become especially critical in recent times due to the high complexity of current data warehousing and mining applications.

In this tutorial, we will conduct a visual exploration of the plan choices made by industrial-strength (commercial and public-domain) optimizers as a function of the input parameter space, whose dimensions include database, query and system-related features. We begin by presenting a suite of diagrams (called plan, cost and cardinality diagrams) that capture the overall behavior of the optimizers over this parameter space. These diagrams are typically remarkably complex and intricate with a large number of plans covering the space, often appearing similar to cubist paintings. They provide a variety of interesting insights, including that current optimizers make extremely fine-grained plan choices, that the plan optimality regions may have highly intricate patterns and irregular boundaries, indicating strongly non-linear cost models; that non-monotonic cost behavior exists where increasing result cardinalities decrease the estimated cost; and, that the basic assumptions underlying the research literature on parametric query optimization often do not hold in practice.

In the next stage, we will show how these complex diagrams can almost always be reduced to much simpler pictures, featuring only a few plans, without materially affecting the query processing quality. The reduction property has several useful implications for the design and usage of query optimizers, including quantifying the redundancy in the plan search space, providing better candidates for plan-cacheing, enhancing the viability of PQO techniques, improving the efficiency of least-expected-cost plans, and minimizing the overheads of adaptive query processing techniques.

15th *International Conference on Management of Data*
COMAD 2009, Mysore, India, December 9–12, 2009

We will present the plan diagram reduction issue from theoretical, statistical and empirical perspectives. Our analysis shows that reduction is an NP-hard problem in general, and remains so even for interesting constrained variations. We will present an online greedy reduction algorithm with tight and optimal performance guarantees, whose complexity scales linearly with the number of plans in the diagram. We then present an estimator that is able to accurately predict the location of the best tradeoff between the query processing quality and the reduction in plan cardinality.

In the last stage, we will focus on the chronic problem of selectivity estimation errors faced by database systems, and demonstrate how our plan diagram reduction scheme can be extended to identify plans that are comparatively robust to such errors. The extension is based on a generalized mathematical characterization of plan cost behavior over the parameter space, that lends itself to efficiently establishing guarantees on the behavior of the substitute plans as compared to the optimizer's standard choices.

Finally, we will show how the diagram post-processing techniques can be incorporated directly into the online query optimization process itself, resulting in a new query optimizer design that delivers a small and select set of robust plans to execute user queries.

All the above-mentioned concepts have been implemented in the (free) Picasso query optimizer visualization tool, and validated on a suite of popular industrial-strength query optimizers with the TPC-H and TPC-DS benchmarks. The Picasso tool is now in use by major database vendors and academic institutions worldwide.

The target audience of the tutorial includes developers, researchers and students interested in database internals, and the pre-requisite is basic knowledge of SQL query processing in relational database engines.

Presenter

Jayant R. Haritsa is on the faculty of the Supercomputer Education & Research Centre and the Department of Computer Science & Automation at the Indian Institute of Science, Bangalore. He received the BTech degree in Electronics and Communications Engineering from the Indian Institute of Technology (Madras), and the MS and PhD degrees in Computer Science from the University of Wisconsin (Madison). His research interests are in database system design and evaluation. Personal web-site: `http://dsl.serc.iisc.ernet.in/~haritsa`.

Primary References

[1] N. Reddy and J. Haritsa, "Analyzing Plan Diagrams of Database Query Optimizers", *Proc. of 31st Intl. Conf. on Very Large Data Bases (VLDB)*, August 2005.

[2] Harish D., P. Darera and J. Haritsa, "On the Production of Anorexic Plan Diagrams", *Proc. of 33rd Intl. Conf. on Very Large Data Bases (VLDB)*, September 2007.

[3] Harish D., P. Darera and J. Haritsa, "Robust Plans through Plan Diagram Reduction", *Proc. of 34th Intl. Conf. on Very Large Data Bases (VLDB)*, August 2008.

[4] A. Dey, S. Bhaumik, Harish D. and J. Haritsa, "Efficiently Approximating Query Optimizer Plan Diagrams", *Proc. of 34th Intl. Conf. on Very Large Data Bases (VLDB)*, August 2008.

[5] M. Abhirama, S. Bhaumik, A. Dey, H. Shrimal and J. Haritsa, "On the Stability of Plan Costs and the Costs of Plan Stability", *(to appear) Proc. of 36th Intl. Conf. on Very Large Data Bases (VLDB)*, September 2010.
Technical report version available at `http://dsl.serc.iisc.ernet.in/publications/report/TR/TR-2009-01.pdf`.

[6] Picasso Project Site: `http://dsl.serc.iisc.ernet.in/projects/PICASSO/picasso.html`

Secondary References

[1] B. Babcock and S. Chaudhuri, "Towards a Robust Query Optimizer: A Principled and Practical Approach", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 2005.

[2] S. Babu, P. Bizarro and D. DeWitt, "Proactive Re-Optimization", *Proc. of ACM Sigmod Intl. Conf. on Management of Data*, June 2005.

[3] P. Bizarro, N. Bruno and D. DeWitt, "Progressive Parametric Query Optimization", *IEEE TKDE*, 21(4), April 2009.

[4] F. Chu, J. Halpern and P. Seshadri, "Least Expected Cost Query Optimization: An Exercise in Utility", *Proc. of ACM Symp. on Principles of Database Systems (PODS)*, May 1999.

[5] F. Chu, J. Halpern and J. Gehrke, "Least Expected Cost Query Optimization: What Can We Expect", *Proc. of ACM Symp. on Principles of Database Systems (PODS)*, May 2002.

[6] A. Deshpande, Z. Ives and V. Raman, "Adaptive Query Processing", *Foundations and Trends in Databases*, 2007.

[7] A. Hulgeri and S. Sudarshan, "Parametric Query Optimization for Linear and Piecewise Linear Cost Functions", *Proc. of 28th Intl. Conf. on Very Large Data Bases (VLDB)*, August 2002.

[8] A. Hulgeri and S. Sudarshan, "AniPQO: Almost Non-intrusive Parametric Query Optimization for Nonlinear Cost Functions", *Proc. of 29th Intl. Conf. on Very Large Data Bases (VLDB)*, September 2003.

[9] Y. Ioannidis and S. Christodoulakis, "On the Propagation of Errors in the Size of Join Results", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1991.

[10] N. Kabra and D. DeWitt, "Efficient Mid-Query Re-Optimization of Sub-Optimal Query Execution Plans", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1998.

[11] L. Mackert and G. Lohman, "R* Optimizer Validation and Performance Evaluation for Local Queries", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1986.

[12] V. Markl, V. Raman, D. Simmen, G. Lohman, H. Pirahesh and M. Cilimdzic, "Robust Query Processing through Progressive Optimization", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 2004.

[13] P. Selinger, M. Astrahan, D. Chamberlin, R. Lorie and T. Price, "Access Path Selection in a Relational Database System", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 1979.

[14] M. Stillger, G. Lohman, V. Markl and M. Kandil, "LEO, DB2's LEarning Optimizer", *Proc. of 27th VLDB Intl. Conf. on Very Large Data Bases (VLDB)*, September 2001.

[15] TPC-H Website: *http://www.tpc.org/tpch*

[16] TPC-DS Website: *http://www.tpc.org/tpcds*