

# CHAPTER 16



## Query Optimization

**Query optimization** is the process of selecting the most efficient query-evaluation plan from among the many strategies usually possible for processing a given query, especially if the query is complex. We do not expect users to write their queries so that they can be processed efficiently. Rather, we expect the system to construct a query-evaluation plan that minimizes the cost of query evaluation. This is where query optimization comes into play.

One aspect of optimization occurs at the relational-algebra level, where the system attempts to find an expression that is equivalent to the given expression, but more efficient to execute. Another aspect is selecting a detailed strategy for processing the query, such as choosing the algorithm to use for executing an operation, choosing the specific indices to use, and so on.

The difference in cost (in terms of evaluation time) between a good strategy and a bad strategy is often substantial and may be several orders of magnitude. Hence, it is worthwhile for the system to spend a substantial amount of time on the selection of a good strategy for processing a query, even if the query is executed only once.

### Bibliographical Notes

The seminal work of [Selinger et al. (1979)] describes access-path selection in the System R optimizer, which was one of the earliest relational-query optimizers. Query processing in Starburst, described in [Haas et al. (1989)], forms the basis for query optimization in IBM DB2.

[Graefe and McKenna (1993)] describe Volcano, an equivalence-rule-based query optimizer that, along with its successor Cascades ([Graefe (1995)]), forms the basis of query optimization in Microsoft SQL Server.

Estimation of statistics of query results, such as result size, is addressed by [Ioannidis and Poosala (1995)], [Poosala et al. (1996)], and [Ganguly et al. (1996)], among others. Nonuniform distributions of values cause problems for estimation of query size and cost. Cost-estimation techniques that use histograms of value distributions have been proposed to tackle the problem. [Ioannidis and Christodoulakis (1993)], [Ioan-

nidis and Poosala (1995)], and [Poosala et al. (1996)] present results in this area. The use of random sampling for constructing histograms is well known in statistics, but issues in histogram construction in the context of databases is discussed in [Chaudhuri et al. (1998)].

[Klug (1982)] was an early work on optimization of relational-algebra expressions with aggregate functions. Optimization of queries with aggregation is addressed by [Yan and Larson (1995)] and [Chaudhuri and Shim (1994)]. Optimization of queries containing outer joins is described in [Rosenthal and Reiner (1984), Galindo-Legaria and Rosenthal (1992)], and [Galindo-Legaria (1994)]. Optimization of top- $K$  queries is addressed in [Carey and Kossmann (1998)] and [Bruno et al. (2002)].

Optimization of nested subqueries is discussed in [Kim (1982), Ganski and Wong (1987), Dayal (1987), Seshadri et al. (1996)] and [Galindo-Legaria and Joshi (2001)].

[Blakeley et al. (1986)] describe techniques for maintenance of materialized views. Optimization of materialized view maintenance plans is described by [Vista (1998)] and [Mistry et al. (2001)]. Query optimization in the presence of materialized views is addressed by [Chaudhuri et al. (1995)]. Index selection and materialized view selection are addressed by [Ross et al. (1996)], and [Chaudhuri and Narasayya (1997)].

Optimization of top- $K$  queries is addressed in [Carey and Kossmann (1998)] and [Bruno et al. (2002)]. A collection of techniques for join minimization has been grouped under the name *tableau optimization*. The notion of a tableau was introduced by [Aho et al. (1979b)] and [Aho et al. (1979a)], and was further extended by [Sagiv and Yannakakis (1981)].

Parametric query-optimization algorithms have been proposed by [Ioannidis et al. (1992)], [Ganguly (1998)] and [Hulgeri and Sudarshan (2003)]. [Sellis (1988)] was an early work on multiquery optimization, while [Roy et al. (2000)] showed how to integrate multi-query optimization into a Volcano-based query optimizer.

[Galindo-Legaria et al. (2004)] describes query processing and optimization for database updates, including optimization of index maintenance, materialized view maintenance plans, and integrity constraint checking, along with techniques to handle the Halloween problem.

## Bibliography

- [Aho et al. (1979a)] A. V. Aho, C. Beeri, and J. D. Ullman, “The Theory of Joins in Relational Databases”, *ACM Transactions on Database Systems*, Volume 4, Number 3 (1979), pages 297–314.
- [Aho et al. (1979b)] A. V. Aho, Y. Sagiv, and J. D. Ullman, “Equivalences among Relational Expressions”, *SIAM Journal of Computing*, Volume 8, Number 2 (1979), pages 218–246.
- [Blakeley et al. (1986)] J. A. Blakeley, P. Larson, and F. W. Tompa, “Efficiently Updating Materialized Views”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (1986), pages 61–71.

- [**Bruno et al. (2002)**] N. Bruno, S. Chaudhuri, and L. Gravano, “Top-k Selection Queries Over Relational Databases: Mapping Strategies and Performance Evaluation”, *ACM Transactions on Database Systems*, Volume 27, Number 2 (2002), pages 153–187.
- [**Carey and Kossmann (1998)**] M. J. Carey and D. Kossmann, “Reducing the Braking Distance of an SQL Query Engine”, In *Proc. of the International Conf. on Very Large Databases (1998)*, pages 158–169.
- [**Chaudhuri and Narasayya (1997)**] S. Chaudhuri and V. Narasayya, “An Efficient Cost-Driven Index Selection Tool for Microsoft SQL Server”, In *Proc. of the International Conf. on Very Large Databases (1997)*, pages 146–155.
- [**Chaudhuri and Shim (1994)**] S. Chaudhuri and K. Shim, “Including Group-By in Query Optimization”, In *Proc. of the International Conf. on Very Large Databases (1994)*, pages 354–366.
- [**Chaudhuri et al. (1995)**] S. Chaudhuri, R. Krishnamurthy, S. Potamianos, and K. Shim, “Optimizing Queries with Materialized Views”, In *Proc. of the International Conf. on Data Engineering (1995)*, pages 190–200.
- [**Chaudhuri et al. (1998)**] S. Chaudhuri, R. Motwani, and V. Narasayya, “Random Sampling for Histogram Construction: How Much is Enough?”, In *Proc. of the ACM SIGMOD Conf. on Management of Data (1998)*, pages 436–447.
- [**Dayal (1987)**] U. Dayal, “Of Nests and Trees: A Unified Approach to Processing Queries that Contain Nested Subqueries, Aggregates and Quantifiers”, In *Proc. of the International Conf. on Very Large Databases (1987)*, pages 197–208.
- [**Galindo-Legaria (1994)**] C. Galindo-Legaria, “Outerjoins as Disjunctions”, In *Proc. of the ACM SIGMOD Conf. on Management of Data (1994)*, pages 348–358.
- [**Galindo-Legaria and Joshi (2001)**] C. A. Galindo-Legaria and M. M. Joshi, “Orthogonal Optimization of Subqueries and Aggregation”, In *Proc. of the ACM SIGMOD Conf. on Management of Data (2001)*, pages 571–581.
- [**Galindo-Legaria and Rosenthal (1992)**] C. Galindo-Legaria and A. Rosenthal, “How to Extend a Conventional Optimizer to Handle One- and Two-Sided Outerjoin”, In *Proc. of the International Conf. on Data Engineering (1992)*, pages 402–409.
- [**Galindo-Legaria et al. (2004)**] C. Galindo-Legaria, S. Stefani, and F. Waas, “Query Processing for SQL Updates”, In *Proc. of the ACM SIGMOD Conf. on Management of Data (2004)*, pages 844–849.
- [**Ganguly (1998)**] S. Ganguly, “Design and Analysis of Parametric Query Optimization Algorithms”, In *Proc. of the International Conf. on Very Large Databases (1998)*, pages 228–238.
- [**Ganguly et al. (1996)**] S. Ganguly, P. Gibbons, Y. Matias, and A. Silberschatz, “A Sampling Algorithm for Estimating Join Size”, In *Proc. of the ACM SIGMOD Conf. on Management of Data (1996)*, pages 271–281.
- [**Ganski and Wong (1987)**] R. A. Ganski and H. K. T. Wong, “Optimization of Nested SQL Queries Revisited”, In *Proc. of the ACM SIGMOD Conf. on Management of Data (1987)*, pages 23–33.

- [Graefe (1995)] G. Graefe, “The Cascades Framework for Query Optimization”, *Data Engineering Bulletin*, Volume 18, Number 3 (1995), pages 19–29.
- [Graefe and McKenna (1993)] G. Graefe and W. McKenna, “The Volcano Optimizer Generator”, In *Proc. of the International Conf. on Data Engineering* (1993), pages 209–218.
- [Haas et al. (1989)] L. M. Haas, J. C. Freytag, G. M. Lohman, and H. Pirahesh, “Extensible Query Processing in Starburst”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (1989), pages 377–388.
- [Hulgeri and Sudarshan (2003)] A. Hulgeri and S. Sudarshan, “AniPQO: Almost Non-Intrusive Parametric Query Optimization for Non-Linear Cost Functions”, In *Proc. of the International Conf. on Very Large Databases* (2003), pages 766–777.
- [Ioannidis and Christodoulakis (1993)] Y. Ioannidis and S. Christodoulakis, “Optimal Histograms for Limiting Worst-Case Error Propagation in the Size of Join Results”, *ACM Transactions on Database Systems*, Volume 18, Number 4 (1993), pages 709–748.
- [Ioannidis and Poosala (1995)] Y. E. Ioannidis and V. Poosala, “Balancing Histogram Optimality and Practicality for Query Result Size Estimation”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (1995), pages 233–244.
- [Ioannidis et al. (1992)] Y. E. Ioannidis, R. T. Ng, K. Shim, and T. K. Sellis, “Parametric Query Optimization”, In *Proc. of the International Conf. on Very Large Databases* (1992), pages 103–114.
- [Kim (1982)] W. Kim, “On Optimizing an SQL-like Nested Query”, *ACM Transactions on Database Systems*, Volume 3, Number 3 (1982), pages 443–469.
- [Klug (1982)] A. Klug, “Equivalence of Relational Algebra and Relational Calculus Query Languages Having Aggregate Functions”, *Journal of the ACM*, Volume 29, Number 3 (1982), pages 699–717.
- [Mistry et al. (2001)] H. Mistry, P. Roy, S. Sudarshan, and K. Ramamritham, “Materialized View Selection and Maintenance Using Multi-Query Optimization”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (2001), pages 307–318.
- [Poosala et al. (1996)] V. Poosala, Y. E. Ioannidis, P. J. Haas, and E. J. Shekita, “Improved Histograms for Selectivity Estimation of Range Predicates”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (1996), pages 294–305.
- [Rosenthal and Reiner (1984)] A. Rosenthal and D. Reiner, “Extending the Algebraic Framework of Query Processing to Handle Outerjoins”, In *Proc. of the International Conf. on Very Large Databases* (1984), pages 334–343.
- [Ross et al. (1996)] K. Ross, D. Srivastava, and S. Sudarshan, “Materialized View Maintenance and Integrity Constraint Checking: Trading Space for Time”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (1996), pages 447–458.
- [Roy et al. (2000)] P. Roy, S. Seshadri, S. Sudarshan, and S. Bhubhe, “Efficient and Extensible Algorithms for Multi-Query Optimization”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (2000), pages 249–260.

- [**Sagiv and Yannakakis (1981)**] Y. Sagiv and M. Yannakakis, “Equivalence among Relational Expressions with the Union and Difference Operators”, *Proc. of the ACM SIGMOD Conf. on Management of Data* (1981), pages 633–655.
- [**Selinger et al. (1979)**] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price, “Access Path Selection in a Relational Database System”, In *Proc. of the ACM SIGMOD Conf. on Management of Data* (1979), pages 23–34.
- [**Sellis (1988)**] T. K. Sellis, “Multiple Query Optimization”, *ACM Transactions on Database Systems*, Volume 13, Number 1 (1988), pages 23–52.
- [**Seshadri et al. (1996)**] P. Seshadri, H. Pirahesh, and T. Y. C. Leung, “Complex Query Decorrelation”, In *Proc. of the International Conf. on Data Engineering* (1996), pages 450–458.
- [**Vista (1998)**] D. Vista, “Integration of Incremental View Maintenance into Query Optimizers”, In *Proc. of the International Conf. on Extending Database Technology* (1998), pages 374–388.
- [**Yan and Larson (1995)**] W. P. Yan and P. A. Larson, “Eager Aggregation and Lazy Aggregation”, In *Proc. of the International Conf. on Very Large Databases* (1995), pages 345–357.

## Credits

The photo of the sailboats in the beginning of the chapter is due to ©Pavel Nesvadba/Shutterstock.

