

# Bibliography

- [1] R. Abbott and H. Garcia-Molina. Scheduling real-time transactions: A performance evaluation. *ACM Transactions on Database Systems*, 17(3), 1992.
- [2] S. Abiteboul. Querying semi-structured data. In *Intl. Conf. on Database Theory*, 1997.
- [3] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [4] S. Abiteboul and P. Kanellakis. Object identity as a query language primitive. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1989.
- [5] S. Abiteboul and V. Vianu. Regular path queries with constraints. In *Proc. ACM Symp. on Principles of Database Systems*, 1997.
- [6] A. Aboulnaga, A. R. Alameldeen, and J. F. Naughton. Estimating the selectivity of XML path expressions for Internet scale applications. In *Proceedings of VLDB*, 2001.
- [7] S. Acharya, P. B. Gibbons, V. Poosala, and S. Ramaswamy. The Aqua approximate query answering system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 574–576. ACM Press, 1999.
- [8] S. Acharya, P. B. Gibbons, V. Poosala, and S. Ramaswamy. Join synopses for approximate query answering. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 275–286. ACM Press, 1999.
- [9] K. Achyutuni, E. Omiecinski, and S. Navathe. Two techniques for on-line index modification in shared nothing parallel databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [10] S. Adali, K. Candan, Y. Papakonstantinou, and V. Subrahmanian. Query caching and optimization in distributed mediator systems. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [11] M. E. Adiba. Derived relations: A unified mechanism for views, snapshots and distributed data. In *Proc. Intl. Conf. on Very Large Databases*, 1981.

- [12] S. Agarwal, R. Agrawal, P. Deshpande, A. Gupta, J. Naughton, R. Ramakrishnan, and S. Sarawagi. On the computation of multidimensional aggregates. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [13] R. C. Agarwal, C. C. Aggarwal, and V. V. V. Prasad. A tree projection algorithm for generation of frequent item sets. *Journal of Parallel and Distributed Computing*, 61(3):350–371, 2001.
- [14] D. Agrawal and A. El Abbadi. The generalized tree quorum protocol: An efficient approach for managing replicated data. *ACM Transactions on Database Systems*, 17(4), 1992.
- [15] D. Agrawal, A. El Abbadi, and R. Jeffers. Using delayed commitment in locking protocols for real-time databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [16] R. Agrawal, M. Carey, and M. Livny. Concurrency control performance-modeling: Alternatives and implications. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1985.
- [17] R. Agrawal and D. DeWitt. Integrated concurrency control and recovery mechanisms: Design and performance evaluation. *ACM Transactions on Database Systems*, 10(4):529–564, 1985.
- [18] R. Agrawal and N. Gehani. ODE (Object Database and Environment): The language and the data model. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1989.
- [19] R. Agrawal, J. E. Gehrke, D. Gunopulos, and P. Raghavan. Automatic subspace clustering of high dimensional data for data mining. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1998.
- [20] R. Agrawal, T. Imielinski, and A. Swami. Database mining: A performance perspective. *IEEE Transactions on Knowledge and Data Engineering*, 5(6):914–925, December 1993.
- [21] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo. Fast discovery of association rules. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, chapter 12, pages 307–328. AAAI/MIT Press, 1996.
- [22] R. Agrawal, G. Psaila, E. Wimmers, and M. Zaot. Querying shapes of histories. In *Proc. Intl. Conf. on Very Large Databases*, 1995.
- [23] R. Agrawal and J. Shafer. Parallel mining of association rules. *IEEE Transactions on Knowledge and Data Engineering*, 8(6):962–969, 1996.
- [24] R. Agrawal and R. Srikant. Mining sequential patterns. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1995.

- [25] R. Agrawal, P. Stolorz, and G. Piatetsky-Shapiro, editors. *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1998.
- [26] R. Ahad, K. BapaRao, and D. McLeod. On estimating the cardinality of the projection of a database relation. *ACM Transactions on Database Systems*, 14(1):28–40, 1989.
- [27] C. Ahlberg and E. Wistrand. IVEE: An information visualization exploration environment. In *Intl. Symp. on Information Visualization*, 1995.
- [28] A. Aho, C. Beeri, and J. Ullman. The theory of joins in relational databases. *ACM Transactions on Database Systems*, 4(3):297–314, 1979.
- [29] A. Aho, J. Hopcroft, and J. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1983.
- [30] A. Aho, Y. Sagiv, and J. Ullman. Equivalences among relational expressions. *SIAM Journal of Computing*, 8(2):218–246, 1979.
- [31] A. Aiken, J. Chen, M. Stonebraker, and A. Woodruff. Tioga-2: A direct manipulation database visualization environment. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1996.
- [32] A. Aiken, J. Widom, and J. Hellerstein. Static analysis techniques for predicting the behavior of active database rules. *ACM Transactions on Database Systems*, 20(1):3–41, 1995.
- [33] A. Ailamaki, D. DeWitt, M. Hill, and M. Skounakis. Weaving relations for cache performance. In *Proc. Intl. Conf. on Very Large Data Bases*, 2001.
- [34] N. Alon, P. B. Gibbons, Y. Matias, and M. Szegedy. Tracking join and self-join sizes in limited storage. In *Proc. ACM Symposium on Principles of Database Systems*, Philadelphia, Pennsylvania, 1999.
- [35] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. In *Proc. of the ACM Symp. on Theory of Computing*, pages 20–29, 1996.
- [36] E. Anwar, L. Maugis, and U. Chakravarthy. A new perspective on rule support for object-oriented databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.
- [37] K. Apt, H. Blair, and A. Walker. Towards a theory of declarative knowledge. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufmann, 1988.
- [38] W. Armstrong. Dependency structures of database relationships. In *Proc. IFIP Congress*, 1974.

- [39] G. Arocena and A. O. Mendelzon. WebOQL: restructuring documents, databases and webs. In *Proc. Intl. Conf. on Data Engineering*, 1988.
- [40] M. Astrahan, M. Blasgen, D. Chamberlin, K. Eswaran, J. Gray, P. Griffiths, W. King, R. Lorie, P. McJones, J. Mehl, G. Putzolu, I. Traiger, B. Wade, and V. Watson. System R: a relational approach to database management. *ACM Transactions on Database Systems*, 1(2):97–137, 1976.
- [41] M. Atkinson, P. Bailey, K. Chisholm, P. Cockshott, and R. Morrison. An approach to persistent programming. In *Readings in Object-Oriented Databases*. eds. S.B. Zdonik and D. Maier, Morgan Kaufmann, 1990.
- [42] M. Atkinson and P. Buneman. Types and persistence in database programming languages. *ACM Computing Surveys*, 19(2):105–190, 1987.
- [43] R. Attar, P. Bernstein, and N. Goodman. Site initialization, recovery, and back-up in a distributed database system. *IEEE Transactions on Software Engineering*, 10(6):645–650, 1983.
- [44] P. Atzeni, L. Cabibbo, and G. Mecca. Isalog: A declarative language for complex objects with hierarchies. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1993.
- [45] P. Atzeni and V. De Antonellis. *Relational Database Theory*. Benjamin-Cummings, 1993.
- [46] P. Atzeni, G. Mecca, and P. Merialdo. To weave the web. In *Proc. Intl. Conf. Very Large Data Bases*, 1997.
- [47] R. Avnur, J. Hellerstein, B. Lo, C. Olston, B. Raman, V. Raman, T. Roth, and K. Wylie. Control: Continuous output and navigation technology with refinement online In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1998.
- [48] R. Avnur and J. M. Hellerstein. Eddies: Continuously adaptive query processing. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 261–272. ACM, 2000.
- [49] B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom. Models and issues in data stream systems. In *Proc. ACM Symp. on Principles of Database Systems*, 2002.
- [50] S. Babu and J. Widom. Continuous queries over data streams. *ACM SIGMOD Record*, 30(3):109–120, 2001.
- [51] D. Badal and G. Popek. Cost and performance analysis of semantic integrity validation methods. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1979.

- [52] A. Badia, D. Van Gucht, and M. Gyssens. Querying with generalized quantifiers. In *Applications of Logic Databases*. ed. R. Ramakrishnan, Kluwer Academic, 1995.
- [53] I. Balbin, G. Port, K. Ramamohanarao, and K. Meenakshi. Efficient bottom-up computation of queries on stratified databases. *Journal of Logic Programming*, 11(3):295–344, 1991.
- [54] I. Balbin and K. Ramamohanarao. A generalization of the differential approach to recursive query evaluation. *Journal of Logic Programming*, 4(3):259–262, 1987.
- [55] F. Bancilhon, C. Delobel, and P. Kanellakis. *Building an Object-Oriented Database System*. Morgan Kaufmann, 1991.
- [56] F. Bancilhon and S. Khoshafian. A calculus for complex objects. *Journal of Computer and System Sciences*, 38(2):326–340, 1989.
- [57] F. Bancilhon, D. Maier, Y. Sagiv, and J. Ullman. Magic sets and other strange ways to implement logic programs. In *ACM Symp. on Principles of Database Systems*, 1986.
- [58] F. Bancilhon and R. Ramakrishnan. An amateur’s introduction to recursive query processing strategies. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1986.
- [59] F. Bancilhon and N. Spyrtos. Update semantics of relational views. *ACM Transactions on Database Systems*, 6(4):557–575, 1981.
- [60] E. Baralis, S. Ceri, and S. Paraboschi. Modularization techniques for active rules design. *ACM Transactions on Database Systems*, 21(1):1–29, 1996.
- [61] D. Barbará, W. DuMouchel, C. Faloutsos, P. J. Haas, J. M. Hellerstein, Y. E. Ioannidis, H. V. Jagadish, T. Johnson, R. T. Ng, V. Poosala, K. A. Ross, and K. C. Sevcik. The New Jersey data reduction report. *Data Engineering Bulletin*, 20(4):3–45, 1997.
- [62] R. Barquin and H. Edelstein. *Planning and Designing the Data Warehouse*. PrenticeHall, 1997.
- [63] C. Batini, S. Ceri, and S. Navathe. *Database Design: An Entity Relationship Approach*. Benjamin/Cummings Publishers, 1992.
- [64] C. Batini, M. Lenzerini, and S. Navathe. A comparative analysis of methodologies for database schema integration. *ACM Computing Surveys*, 18(4):323–364, 1986.

- [65] D. Batory, J. Barnett, J. Garza, K. Smith, K. Tsukuda, B. Twichell, and T. Wise. GENESIS: An extensible database management system. In S. Zdonik and D. Maier, editors, *Readings in Object-Oriented Databases*. Morgan Kaufmann, 1990.
- [66] B. Baugsto and J. Greipsland. Parallel sorting methods for large data volumes on a hypercube database computer. In *Proc. Intl. Workshop on Database Machines*, 1989.
- [67] R. J. Bayardo. Efficiently mining long patterns from databases. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, pages 85–93. ACM Press, 1998.
- [68] R. J. Bayardo, R. Agrawal, and D. Gunopulos. Constraint-based rule mining in large, dense databases. *Data Mining and Knowledge Discovery*, 4(2/3):217–240, 2000.
- [69] R. Bayer and E. McCreight. Organization and maintenance of large ordered indexes. *Acta Informatica*, 1(3):173–189, 1972.
- [70] R. Bayer and M. Schkolnick. Concurrency of operations on B-trees. *Acta Informatica*, 9(1):1–21, 1977.
- [71] M. Beck, D. Bitton, and W. Wilkinson. Sorting large files on a backend multiprocessor. *IEEE Transactions on Computers*, 37(7):769–778, 1988.
- [72] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger. The R\* tree: An efficient and robust access method for points and rectangles. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [73] C. Beeri, R. Fagin, and J. Howard. A complete axiomatization of functional and multivalued dependencies in database relations. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1977.
- [74] C. Beeri and P. Honeyman. Preserving functional dependencies. *SIAM Journal of Computing*, 10(3):647–656, 1982.
- [75] C. Beeri and T. Milo. A model for active object-oriented database. In *Proc. Intl. Conf. on Very Large Databases*, 1991.
- [76] C. Beeri, S. Naqvi, R. Ramakrishnan, O. Shmueli, and S. Tsur. Sets and negation in a logic database language (LDL1 ). In *ACM Symp. on Principles of Database Systems*, 1987.
- [77] C. Beeri and R. Ramakrishnan. On the power of magic. In *ACM Symp. on Principles of Database Systems*, 1987.
- [78] D. Bell and J. Grimson. *Distributed Database Systems*. Addison-Wesley, 1992.

- [79] J. Bentley and J. Friedman. Data structures for range searching. *ACM Computing Surveys*, 13(3):397–409, 1979.
- [80] S. Berchtold, C. Bohm, and H.-P. Kriegel. The pyramid-tree: breaking the curse of dimensionality. In *ACM SIGMOD Conf. on the Management of Data*, 1998.
- [81] P. Bernstein. Synthesizing third normal form relations from functional dependencies. *ACM Transactions on Database Systems*, 1(4):277–298, 1976.
- [82] P. Bernstein, B. Blaustein, and E. Clarke. Fast maintenance of semantic integrity assertions using redundant aggregate data. In *Proc. Intl. Conf. on Very Large Databases*, 1980.
- [83] P. Bernstein and D. Chiu. Using semi-joins to solve relational queries. *Journal of the ACM*, 28(1):25–40, 1981.
- [84] P. Bernstein and N. Goodman. Timestamp-based algorithms for concurrency control in distributed database systems. In *Proc. Intl. Conf. on Very Large Databases*, 1980.
- [85] P. Bernstein and N. Goodman. Concurrency control in distributed database systems. *ACM Computing Surveys*, 13(2):185–222, 1981.
- [86] P. Bernstein and N. Goodman. Power of natural semijoins. *SIAM Journal of Computing*, 10(4):751–771, 1981.
- [87] P. Bernstein and N. Goodman. Multiversion concurrency control—Theory and algorithms. *ACM Transactions on Database Systems*, 8(4):465–483, 1983.
- [88] P. Bernstein, N. Goodman, E. Wong, C. Reeve, and J. Rothnie. Query processing in a system for distributed databases (SDD-1 ). *ACM Transactions on Database Systems*, 6(4):602–625, 1981.
- [89] P. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987.
- [90] P. Bernstein and E. Newcomer. *Principles of Transaction Processing*. Morgan Kaufmann, 1997.
- [91] P. Bernstein, D. Shipman, and J. Rothnie. Concurrency control in a system for distributed databases (SDD-1 ). *ACM Transactions on Database Systems*, 5(1):18–51, 1980.
- [92] P. Bernstein, D. Shipman, and W. Wong. Formal aspects of serializability in database concurrency control. *IEEE Transactions on Software Engineering*, 5(3):203–216, 1979.

- [93] K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft. When is nearest neighbor meaningful? In *IEEE International Conference on Database Theory*, 1999.
- [94] K. Beyer and R. Ramakrishnan. Bottom-up computation of sparse and iceberg cubes In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1999.
- [95] B. Bhargava, editor. *Concurrency Control and Reliability in Distributed Systems*. Van Nostrand Reinhold, 1987.
- [96] A. Biliris. The performance of three database storage structures for managing large objects. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [97] J. Biskup and B. Convent. A formal view integration method. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1986.
- [98] J. Biskup, U. Dayal, and P. Bernstein. Synthesizing independent database schemas. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1979.
- [99] D. Bitton and D. DeWitt. Duplicate record elimination in large data files. *ACM Transactions on Database Systems*, 8(2):255–265, 1983.
- [100] J. Blakeley, P.-A. Larson, and F. Tompa. Efficiently updating materialized views. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1986.
- [101] M. Blasgen and K. Eswaran. On the evaluation of queries in a database system. Technical report, IBM FJ (RJ1745), San Jose, 1975.
- [102] P. Bohannon, D. Leinbaugh, R. Rastogi, S. Seshadri, A. Silberschatz, and S. Sudarshan. Logical and physical versioning in main memory databases. In *Proc. Intl. Conf. on Very Large Databases*, 1997.
- [103] P. Bohannon, J. Freire, P. Roy, and J. Simeon. From XML schema to relations: A cost-based approach to XML storage. In *Proceedings of ICDE*, 2002.
- [104] P. Bonnet and D. E. Shasha. *Database Tuning: Principles, Experiments, and Troubleshooting Techniques*. Morgan Kaufmann Publishers, 2002.
- [105] G. Booch, I. Jacobson, and J. Rumbaugh. *The Unified Modeling Language User Guide*. Addison-Wesley, 1998.
- [106] A. Borodin, G. Roberts, J. Rosenthal, and P. Tsaparas. Finding authorities and hubs from link structures on Roberts G.O. the world wide web. In *World Wide Web Conference*, pages 415–429, 2001.
- [107] R. Boyce and D. Chamberlin. SEQUEL: A structured English query language. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1974.

- [108] P. S. Bradley and U. M. Fayyad. Refining initial points for K-Means clustering. In *Proc. Intl. Conf. on Machine Learning*, pages 91–99. Morgan Kaufmann, San Francisco, CA, 1998.
- [109] P. S. Bradley, U. M. Fayyad, and C. Reina. Scaling clustering algorithms to large databases. In *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*, 1998.
- [110] K. Bratbergsengen. Hashing methods and relational algebra operations. In *Proc. Intl. Conf. on Very Large Databases*, 1984.
- [111] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, Belmont. CA, 1984.
- [112] Y. Breitbart, H. Garcia-Molina, and A. Silberschatz. Overview of multidatabase transaction management. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [113] Y. Breitbart, A. Silberschatz, and G. Thompson. Reliable transaction management in a multidatabase system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [114] Y. Breitbart, A. Silberschatz, and G. Thompson. An approach to recovery management in a multidatabase system. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [115] S. Brin, R. Motwani, and C. Silverstein. Beyond market baskets: Generalizing association rules to correlations. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [116] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *Proceedings of 7th World Wide Web Conference*, 1998.
- [117] S. Brin, R. Motwani, J. D. Ullman, and S. Tsur. Dynamic itemset counting and implication rules for market basket data. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, pages 255–264. ACM Press, 1997.
- [118] T. Brinkhoff, H.-P. Kriegel, and R. Schneider. Comparison of approximations of complex objects used for approximation-based query processing in spatial database systems. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1993.
- [119] K. Brown, M. Carey, and M. Livny. Goal-oriented buffer management revisited. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [120] N. Bruno, S. Chaudhuri, and L. Gravano. Top-k selection queries over relational databases: Mapping strategies and performance evaluation. *ACM Transactions on Database Systems*, To appear, 2002.

- [121] F. Bry. Towards an efficient evaluation of general queries: Quantifier and disjunction processing revisited. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1989.
- [122] F. Bry and R. Manthey. Checking consistency of database constraints: A logical basis. In *Proc. Intl. Conf. on Very Large Databases*, 1986.
- [123] P. Buneman and E. Clemons. Efficiently monitoring relational databases. *ACM Transactions on Database Systems*, 4(3), 1979.
- [124] P. Buneman, S. Davidson, G. Hillebrand, and D. Suciu. A query language and optimization techniques for unstructured data. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1996.
- [125] P. Buneman, S. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *Theoretical Computer Science*, 149(1):3–48, 1995.
- [126] D. Burdick, M. Calimlim, and J. E. Gehrke. Mafia: A maximal frequent itemset algorithm for transactional databases. In *Proc. Intl. Conf. on Data Engineering (ICDE)*. IEEE Computer Society, 2001.
- [127] M. Carey. Granularity hierarchies in concurrency control. In *ACM Symp. on Principles of Database Systems*, 1983.
- [128] M. Carey, D. Chamberlin, S. Narayanan, B. Vance, D. Doole, S. Rielau, R. Swagerman, and N. Mattos. O-O, what’s happening to DB2? In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1999.
- [129] M. Carey, D. DeWitt, M. Franklin, N. Hall, M. McAuliffe, J. Naughton, D. Schuh, M. Solomon, C. Tan, O. Tsatalos, S. White, and M. Zwillig. Shoring up persistent applications. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.
- [130] M. Carey, D. DeWitt, G. Graefe, D. Haight, J. Richardson, D. Schuh, E. Shekita, and S. Vandenberg. The EXODUS Extensible DBMS project: An overview. In S. Zdonik and D. Maier, editors, *Readings in Object-Oriented Databases*. Morgan Kaufmann, 1990.
- [131] M. Carey, D. DeWitt, and J. Naughton. The 007 benchmark. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.
- [132] M. Carey, D. DeWitt, J. Naughton, M. Asgarian, J. Gehrke, and D. Shah. The BUCKY object-relational benchmark. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [133] M. Carey, D. DeWitt, J. Richardson, and E. Shekita. Object and file management in the Exodus extensible database system. In *Proc. Intl. Conf. on Very Large Databases*, 1986.

- [134] M. Carey, D. Florescu, Z. Ives, Y. Lu, J. Shanmugasundaram, E. Shekita, and S. Subramanian. XPERANTO: publishing object-relational data as XML. In *Proceedings of the Third International Workshop on the Web and Databases*, May 2000.
- [135] M. Carey and D. Kossman. On saying “Enough Already!” in SQL. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [136] M. Carey and D. Kossman. Reducing the braking distance of an SQL query engine. In *Proc. Intl. Conf. on Very Large Databases*, 1998.
- [137] M. Carey and M. Livny. Conflict detection tradeoffs for replicated data. *ACM Transactions on Database Systems*, 16(4), 1991.
- [138] M. Casanova, L. Tucheran, and A. Furtado. Enforcing inclusion dependencies and referential integrity. In *Proc. Intl. Conf. on Very Large Databases*, 1988.
- [139] M. Casanova and M. Vidal. Towards a sound view integration methodology. In *ACM Symp. on Principles of Database Systems*, 1983.
- [140] S. Castano, M. Fugini, G. Martella, and P. Samarati. *Database Security*. Addison-Wesley, 1995.
- [141] R. Cattell. *The Object Database Standard: ODMG-93 (Release 1.1)*. Morgan Kaufmann, 1994.
- [142] S. Ceri, P. Fraternali, S. Paraboschi, and L. Tanca. Active rule management in Chimera. In J. Widom and S. Ceri, editors, *Active Database Systems*. Morgan Kaufmann, 1996.
- [143] S. Ceri, G. Gottlob, and L. Tanca. *Logic Programming and Databases*. Springer Verlag, 1990.
- [144] S. Ceri and G. Pelagatti. *Distributed Database Design: Principles and Systems*. McGraw-Hill, 1984.
- [145] S. Ceri and J. Widom. Deriving production rules for constraint maintenance. In *Proc. Intl. Conf. on Very Large Databases*, 1990.
- [146] F. Cesarini, M. Missikoff, and G. Soda. An expert system approach for database application tuning. *Data and Knowledge Engineering*, 8:35–55, 1992.
- [147] U. Chakravarthy. Architectures and monitoring techniques for active databases: An evaluation. *Data and Knowledge Engineering*, 16(1):1–26, 1995.
- [148] U. Chakravarthy, J. Grant, and J. Minker. Logic-based approach to semantic query optimization. *ACM Transactions on Database Systems*, 15(2):162–207, 1990.

- [149] D. Chamberlin. *Using the New DB2*. Morgan Kaufmann, 1996.
- [150] D. Chamberlin, M. Astrahan, M. Blasgen, J. Gray, W. King, B. Lindsay, R. Lorie, J. Mehl, T. Price, P. Selinger, M. Schkolnick, D. Slutz, I. Traiger, B. Wade, and R. Yost. A history and evaluation of System R. *Communications of the ACM*, 24(10):632–646, 1981.
- [151] D. Chamberlin, M. Astrahan, K. Eswaran, P. Griffiths, R. Lorie, J. Mehl, P. Reisner, and B. Wade. Sequel 2: a unified approach to data definition, manipulation, and control. *IBM Journal of Research and Development*, 20(6):560–575, 1976.
- [152] D. Chamberlin, D. Florescu, and J. Robie. Quilt: an XML query language for heterogeneous data sources. In *Proceedings of WebDB*, Dallas, TX, May 2000.
- [153] D. Chamberlin, D. Florescu, J. Robie, J. Simeon, and M. Stefanescu. XQuery: A query language for XML. World Wide Web Consortium, <http://www.w3.org/TR/xquery>, Feb 2000.
- [154] A. Chandra and D. Harel. Structure and complexity of relational queries. *J. Computer and System Sciences*, 25:99–128, 1982.
- [155] A. Chandra and P. Merlin. Optimal implementation of conjunctive queries in relational databases. In *Proc. ACM SIGACT Symp. on Theory of Computing*, 1977.
- [156] M. Chandy, L. Haas, and J. Misra. Distributed deadlock detection. *ACM Transactions on Computer Systems*, 1(3):144–156, 1983.
- [157] C. Chang and D. Leu. Multi-key sorting as a file organization scheme when queries are not equally likely. In *Proc. Intl. Symp. on Database Systems for Advanced Applications*, 1989.
- [158] D. Chang and D. Harkey. *Client/ server data access with Java and XML*. John Wiley and Sons, 1998.
- [159] M. Charikar, S. Chaudhuri, R. Motwani, and V. R. Narasayya. Towards estimation error guarantees for distinct values. In *Proc. ACM Symposium on Principles of Database Systems*, pages 268–279. ACM, 2000.
- [160] D. Chatziantoniou and K. Ross. Groupwise processing of relational queries. In *Proc. Intl. Conf. on Very Large Databases*, 1997.
- [161] S. Chaudhuri and U. Dayal. An overview of data warehousing and OLAP technology. *SIGMOD Record*, 26(1):65–74, 1997.
- [162] S. Chaudhuri and D. Madigan, editors. *Proc. ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*. ACM Press, 1999.

- [163] S. Chaudhuri and V. Narasayya. An efficient cost-driven index selection tool for Microsoft SQL Server. In *Proc. Intl. Conf. on Very Large Databases*, 1997.
- [164] S. Chaudhuri and V. R. Narasayya. Autoadmin 'what-if' index analysis utility. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 1998.
- [165] S. Chaudhuri and K. Shim. Optimization of queries with user-defined predicates. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [166] S. Chaudhuri and K. Shim. Optimization queries with aggregate views. In *Intl. Conf. on Extending Database Technology*, 1996.
- [167] S. Chaudhuri, G. Das, and V. R. Narasayya. A robust, optimization-based approach for approximate answering of aggregate queries. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 2001.
- [168] J. Cheiney, P. Faudemay, R. Michel, and J. Thevenin. A reliable parallel backend using multiattribute clustering and select-join operator. In *Proc. Intl. Conf. on Very Large Databases*, 1986.
- [169] C. Chen and N. Roussopoulos. Adaptive database buffer management using query feedback. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [170] C. Chen and N. Roussopoulos. Adaptive selectivity estimation using query feedback. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.
- [171] P. M. Chen, E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson. RAID: High-performance, reliable secondary storage. *ACM Computing Surveys*, 26(2):145–185, June 1994.
- [172] P. P. Chen. The entity-relationship model—toward a unified view of data. *ACM Transactions on Database Systems*, 1(1):9–36, 1976.
- [173] Y. Chen, G. Dong, J. Han, B. W. Wah, and J. Wang. Multi-dimensional regression analysis of time-series data streams. In *Proc. Intl. Conf. on Very Large Data Bases*, 2002.
- [174] D. W. Cheung, J. Han, V. T. Ng, and C. Y. Wong. Maintenance of discovered association rules in large databases: An incremental updating technique. In *Proc. Intl. Conf. Data Engineering*, 1996.
- [175] D. W. Cheung, V. T. Ng, and B. W. Tam. Maintenance of discovered knowledge: A case in multi-level association rules. In *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1996.
- [176] D. Childs. Feasibility of a set theoretical data structure—A general structure based on a reconstructed definition of relation. *Proc. Tri-annual IFIP Conference*, 1968.

- [177] D. Chimenti, R. Gamboa, R. Krishnamurthy, S. Naqvi, S. Tsur, and C. Zaniolo. The ldl system prototype. *IEEE Transactions on Knowledge and Data Engineering*, 2(1):76–90, 1990.
- [178] F. Chin and G. Ozsoyoglu. Statistical database design. *ACM Transactions on Database Systems*, 6(1):113–139, 1981.
- [179] T.-C. Chiueh and L. Huang. Efficient real-time index updates in text retrieval systems.
- [180] J. Chomicki. Real-time integrity constraints. In *ACM Symp. on Principles of Database Systems*, 1992.
- [181] H.-T. Chou and D. DeWitt. An evaluation of buffer management strategies for relational database systems. In *Proc. Intl. Conf. on Very Large Databases*, 1985.
- [182] P. Chrysanthis and K. Ramamritham. Acta: A framework for specifying and reasoning about transaction structure and behavior. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [183] F. Chu, J. Halpern, and P. Seshadri. Least expected cost query optimization: An exercise in utility *ACM Symp. on Principles of Database Systems*, 1999.
- [184] F. Civelek, A. Dogac, and S. Spaccapietra. An expert system approach to view definition and integration. In *Proc. Entity-Relationship Conference*, 1988.
- [185] R. Cochrane, H. Pirahesh, and N. Mattos. Integrating triggers and declarative constraints in SQL database systems. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [186] CODASYL. *Report of the CODASYL Data Base Task Group*. ACM, 1971.
- [187] E. Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387, 1970.
- [188] E. Codd. Further normalization of the data base relational model. In R. Rustin, editor, *Data Base Systems*. Prentice Hall, 1972.
- [189] E. Codd. Relational completeness of data base sub-languages. In R. Rustin, editor, *Data Base Systems*. Prentice Hall, 1972.
- [190] E. Codd. Extending the database relational model to capture more meaning. *ACM Transactions on Database Systems*, 4(4):397–434, 1979.
- [191] E. Codd. Twelve rules for on-line analytic processing. *Computerworld*, April 13 1995.

- [192] L. Colby, T. Griffin, L. Libkin, I. Mumick, and H. Trickey. Algorithms for deferred view maintenance. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [193] L. Colby, A. Kawaguchi, D. Lieuwen, I. Mumick, and K. Ross. Supporting multiple view maintenance policies: Concepts, algorithms, and performance analysis. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [194] D. Comer. The ubiquitous B-tree. *ACM C. Surveys*, 11(2):121–137, 1979.
- [195] D. Connolly, editor. *XML Principles, Tools and Techniques*. O'Reilly & Associates, Sebastopol, USA, 1997.
- [196] B. Cooper, N. Sample, M. J. Franklin, G. R. Hjaltason, and M. Shadmon. A fast index for semistructured data. In *Proceedings of VLDB*, 2001.
- [197] D. Copeland and D. Maier. Making SMALLTALK a database system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1984.
- [198] G. Cornell and K. Abdali. *CGI Programming With Java*. PrenticeHall, 1998.
- [199] C. Cortes, K. Fisher, D. Pregibon, and A. Rogers. Hancock: a language for extracting signatures from data streams. In *Proc. ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*, pages 9–17. AAAI Press, 2000.
- [200] J. Daemen and V. Rijmen. *The Design of Rijndael: AES –The Advanced Encryption Standard (Information Security and Cryptography)*. Springer Verlag, 2002.
- [201] M. Datar, A. Gionis, P. Indyk, and R. Motwani. Maintaining stream statistics over sliding windows. In *Proc. of the Annual ACM-SIAM Symp. on Discrete Algorithms*, 2002.
- [202] C. Date. A critique of the SQL database language. *ACM SIGMOD Record*, 14(3):8–54, 1984.
- [203] C. Date. *Relational Database: Selected Writings*. Addison-Wesley, 1986.
- [204] C. Date. *An Introduction to Database Systems*. Addison-Wesley, 7 edition, 1999.
- [205] C. Date and R. Fagin. Simple conditions for guaranteeing higher normal forms in relational databases. *ACM Transactions on Database Systems*, 17(3), 1992.
- [206] C. Date and D. McGoveran. *A Guide to Sybase and SQL Server*. Addison-Wesley, 1993.

- [207] U. Dayal and P. Bernstein. On the updatability of relational views. In *Proc. Intl. Conf. on Very Large Databases*, 1978.
- [208] U. Dayal and P. Bernstein. On the correct translation of update operations on relational views. *ACM Transactions on Database Systems*, 7(3), 1982.
- [209] P. DeBra and J. Paredaens. Horizontal decompositions for handling exceptions to FDs. In H. Gallaire, J. Minker, and J.-M. Nicolas, editors, *Advances in Database Theory*,. Plenum Press, 1981.
- [210] J. Deep and P. Holfelder. *Developing CGI applications with Perl*. Wiley, 1996.
- [211] C. Delobel. Normalization and hierarchial dependencies in the relational data model. *ACM Transactions on Database Systems*, 3(3):201–222, 1978.
- [212] D. Denning. Secure statistical databases with random sample queries. *ACM Transactions on Database Systems*, 5(3):291–315, 1980.
- [213] D. E. Denning. *Cryptography and Data Security*. Addison-Wesley, 1982.
- [214] M. Derr, S. Morishita, and G. Phipps. The glue-nail deductive database system: Design, implementation, and evaluation. *VLDB Journal*, 3(2):123–160, 1994.
- [215] A. Deshpande. An implementation for nested relational databases. Technical report, PhD thesis, Indiana University, 1989.
- [216] P. Deshpande, K. Ramasamy, A. Shukla, and J. F. Naughton. Caching multidimensional queries using chunks. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 1998.
- [217] A. Deutsch, M. Fernandez, D. Florescu, A. Levy, and D. Suciu. XML-QL: A query language for XML. World Wide Web Consortium, <http://www.w3.org/TR/NOTE-xml-q1>, Aug 1998.
- [218] O. e. a. Deux. The story of O2. *IEEE Transactions on Knowledge and Data Engineering*, 2(1), 1990.
- [219] D. DeWitt, H.-T. Chou, R. Katz, and A. Klug. Design and implementation of the Wisconsin Storage System. *Software Practice and Experience*, 15(10):943–962, 1985.
- [220] D. DeWitt, R. Gerber, G. Graefe, M. Heytens, K. Kumar, and M. Muralikrishna. Gamma—A high performance dataflow database machine. In *Proc. Intl. Conf. on Very Large Databases*, 1986.
- [221] D. DeWitt and J. Gray. Parallel database systems: The future of high-performance database systems. *Communications of the ACM*, 35(6):85–98, 1992.

- [222] D. DeWitt, R. Katz, F. Olken, L. Shapiro, M. Stonebraker, and D. Wood. Implementation techniques for main memory databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1984.
- [223] D. DeWitt, J. Naughton, and D. Schneider. Parallel sorting on a shared-nothing architecture using probabilistic splitting. In *Proc. Conf. on Parallel and Distributed Information Systems*, 1991.
- [224] D. DeWitt, J. Naughton, D. Schneider, and S. Seshadri. Practical skew handling in parallel joins. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [225] O. Diaz, N. Paton, and P. Gray. Rule management in object-oriented databases: A uniform approach. In *Proc. Intl. Conf. on Very Large Databases*, 1991.
- [226] S. Dietrich. Extension tables: Memo relations in logic programming. In *Proc. Intl. Symp. on Logic Programming*, 1987.
- [227] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [228] P. Domingos and G. Hulten. Mining high-speed data streams. In *Proc. ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*. AAAI Press, 2000.
- [229] D. Donjerkovic and R. Ramakrishnan. Probabilistic optimization of top N queries In *Proc. Intl. Conf. on Very Large Databases*, 1999.
- [230] W. Du and A. Elmagarmid. Quasi-serializability: A correctness criterion for global concurrency control in interbase. In *Proc. Intl. Conf. on Very Large Databases*, 1989.
- [231] W. Du, R. Krishnamurthy, and M.-C. Shan. Query optimization in a heterogeneous DBMS. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [232] R. C. Dubes and A. Jain. *Clustering Methodologies in Exploratory Data Analysis, Advances in Computers*. Academic Press, New York, 1980.
- [233] N. Duppel. Parallel SQL on TANDEM 's NonStop SQL. *IEEE COMP-CON*, 1989.
- [234] H. Edelstein. The challenge of replication, Parts 1 and 2. *DBMS: Database and Client-Server Solutions*, 1995.
- [235] W. Effelsberg and T. Haerder. Principles of database buffer management. *ACM Transactions on Database Systems*, 9(4):560–595, 1984.
- [236] M. H. Eich. A classification and comparison of main memory database recovery techniques. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1987.

- [237] A. Eisenberg and J. Melton. SQL:1999 , formerly known as SQL 3 *ACM SIGMOD Record*, 28(1):131–138, 1999.
- [238] A. El Abbadi. Adaptive protocols for managing replicated distributed databases. In *IEEE Symp. on Parallel and Distributed Processing*, 1991.
- [239] A. El Abbadi, D. Skeen, and F. Cristian. An efficient, fault-tolerant protocol for replicated data management. In *ACM Symp. on Principles of Database Systems*, 1985.
- [240] C. Ellis. Concurrency in Linear Hashing. *ACM Transactions on Database Systems*, 12(2):195–217, 1987.
- [241] A. Elmagarmid. *Database Transaction Models for Advanced Applications*. Morgan Kaufmann, 1992.
- [242] A. Elmagarmid, J. Jing, W. Kim, O. Bukhres, and A. Zhang. Global commitability in multidatabase systems. *IEEE Transactions on Knowledge and Data Engineering*, 8(5):816–824, 1996.
- [243] A. Elmagarmid, A. Sheth, and M. Liu. Deadlock detection algorithms in distributed database systems. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1986.
- [244] R. Elmasri and S. Navathe. Object integration in database design. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1984.
- [245] R. Elmasri and S. Navathe. *Fundamentals of Database Systems*. Benjamin-Cummings, 3 edition, 2000.
- [246] R. Epstein. Techniques for processing of aggregates in relational database systems. Technical report, UC-Berkeley, Electronics Research Laboratory, M798, 1979.
- [247] R. Epstein, M. Stonebraker, and E. Wong. Distributed query processing in a relational data base system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1978.
- [248] M. Ester, H.-P. Kriegel, J. Sander, M. Wimmer, and X. Xu. Incremental clustering for mining in a data warehousing environment. In *Proc. Intl. Conf. On Very Large Data Bases*, 1998.
- [249] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc. Intl. Conf. on Knowledge Discovery in Databases and Data Mining*, 1995.
- [250] M. Ester, H.-P. Kriegel, and X. Xu. A database interface for clustering in large spatial databases. In *Proc. Intl. Conf. on Knowledge Discovery in Databases and Data Mining*, 1995.

- [251] K. Eswaran and D. Chamberlin. Functional specification of a subsystem for data base integrity. In *Proc. Intl. Conf. on Very Large Databases*, 1975.
- [252] K. Eswaran, J. Gray, R. Lorie, and I. Traiger. The notions of consistency and predicate locks in a data base system. *Communications of the ACM*, 19(11):624–633, 1976.
- [253] R. Fagin. Multivalued dependencies and a new normal form for relational databases. *ACM Transactions on Database Systems*, 2(3):262–278, 1977.
- [254] R. Fagin. Normal forms and relational database operators. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1979.
- [255] R. Fagin. A normal form for relational databases that is based on domains and keys. *ACM Transactions on Database Systems*, 6(3):387–415, 1981.
- [256] R. Fagin, J. Nievergelt, N. Pippenger, and H. Strong. Extendible Hashing—a fast access method for dynamic files. *ACM Transactions on Database Systems*, 4(3), 1979.
- [257] C. Faloutsos. Access methods for text. *ACM Computing Surveys*, 17(1):49–74, 1985.
- [258] C. Faloutsos. *Searching Multimedia Databases by Content* Kluwer Academic, 1996.
- [259] C. Faloutsos and S. Christodoulakis. Signature files: An access method for documents and its analytical performance evaluation. *ACM Transactions on Office Information Systems*, 2(4):267–288, 1984.
- [260] C. Faloutsos and H. Jagadish. On B-Tree indices for skewed distributions. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [261] C. Faloutsos, R. Ng, and T. Sellis. Predictive load control for flexible buffer allocation. In *Proc. Intl. Conf. on Very Large Databases*, 1991.
- [262] C. Faloutsos, M. Ranganathan, and Y. Manolopoulos. Fast subsequence matching in time-series databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.
- [263] C. Faloutsos and S. Roseman. Fractals for secondary key retrieval. In *ACM Symp. on Principles of Database Systems*, 1989.
- [264] M. Fang, N. Shivakumar, H. Garcia-Molina, R. Motwani, and J. D. Ullman. Computing iceberg queries efficiently. In *Proc. Intl. Conf. On Very Large Data Bases*, 1998.
- [265] U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. The KDD process for extracting useful knowledge from volumes of data. *Communications of the ACM*, 39(11):27–34, 1996.

- [266] U. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors. *Advances in Knowledge Discovery and Data Mining*. MIT Press, 1996.
- [267] U. Fayyad and E. Simoudis. Data mining and knowledge discovery: Tutorial notes. In *Intl. Joint Conf. on Artificial Intelligence*, 1997.
- [268] U. M. Fayyad and R. Uthurusamy, editors. *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1995.
- [269] M. Fernandez, D. Florescu, J. Kang, A. Y. Levy, and D. Suciu. STRUDEL: A Web site management system. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1997.
- [270] M. Fernandez, D. Florescu, A. Y. Levy, and D. Suciu. A query language for a Web -site management system. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 26(3):4–11, 1997.
- [271] M. Fernandez, D. Suciu, and W. Tan. SilkRoute: trading between relations and XML. In *Proceedings of the WWW9*, 2000.
- [272] S. Finkelstein, M. Schkolnick, and P. Tiberio. Physical database design for relational databases. *IBM Research Review RJ5034*, 1986.
- [273] D. Fishman, D. Beech, H. Cate, E. Chow, T. Connors, J. Davis, N. Derrett, C. Hoch, W. Kent, P. Lyngbaek, B. Mahbod, M.-A. Neimat, T. Ryan, and M.-C. Shan. Iris: an object-oriented database management system *ACM Transactions on Office Information Systems*, 5(1):48–69, 1987.
- [274] C. Fleming and B. von Halle. *Handbook of Relational Database Design*. Addison-Wesley, 1989.
- [275] D. Florescu, A. Y. Levy, and A. O. Mendelzon. Database techniques for the World-Wide Web: A survey. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 27(3):59–74, 1998.
- [276] W. Ford and M. S. Baum. *Secure Electronic Commerce: Building the Infrastructure for Digital Signatures and Encryption (2nd Edition)*. Prentice Hall, 2000.
- [277] F. Fotouhi and S. Pramanik. Optimal secondary storage access sequence for performing relational join. *IEEE Transactions on Knowledge and Data Engineering*, 1(3):318–328, 1989.
- [278] M. Fowler and K. Scott. *UML Distilled: Applying the Standard Object Modeling Language*. Addison-Wesley, 1999.
- [279] W. B. Frakes and R. Baeza-Yates, editors. *Information Retrieval: Data Structures and Algorithms*. PrenticeHall, 1992.

- [280] P. Franaszek, J. Robinson, and A. Thomasian. Concurrency control for high contention environments. *ACM Transactions on Database Systems*, 17(2), 1992.
- [281] P. Franaszek, J. Robinson, and A. Thomasian. Access invariance and its use in high contention environments. In *Proc. IEEE International Conference on Data Engineering*, 1990.
- [282] M. Franklin. Concurrency control and recovery. In *Handbook of Computer Science, A.B. Tucker (ed.)*, CRC Press, 1996.
- [283] M. Franklin, M. Carey, and M. Livny. Local disk caching for client-server database systems. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [284] M. Franklin, B. Jonsson, and D. Kossman. Performance tradeoffs for client-server query processing. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [285] P. Fraternali and L. Tanca. A structured approach for the definition of the semantics of active databases. *ACM Transactions on Database Systems*, 20(4):414–471, 1995.
- [286] M. W. Freeston. The BANG file: A new kind of Grid File. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [287] J. Freytag. A rule-based view of query optimization. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [288] O. Friesen, A. Lefebvre, and L. Vieille. VALIDITY: Applications of a DOOD system. In *Intl. Conf. on Extending Database Technology*, 1996.
- [289] J. Fry and E. Sibley. Evolution of data-base management systems. *ACM Computing Surveys*, 8(1):7–42, 1976.
- [290] N. Fuhr. A decision-theoretic approach to database selection in networked ir. *ACM Transactions on Database Systems*, 17(3), 1999.
- [291] T. Fukuda, Y. Morimoto, S. Morishita, and T. Tokuyama. Mining optimized association rules for numeric attributes. In *ACM Symp. on Principles of Database Systems*, 1996.
- [292] A. Furtado and M. Casanova. Updating relational views. In *Query Processing in Database Systems*. eds. W. Kim, D.S. Reiner and D.S. Batory, Springer-Verlag, 1985.
- [293] S. Fushimi, M. Kitsuregawa, and H. Tanaka. An overview of the systems software of a parallel relational database machine: Grace. In *Proc. Intl. Conf. on Very Large Databases*, 1986.
- [294] V. Gaede and O. Guenther. Multidimensional access methods. *Computing Surveys*, 30(2):170–231, 1998.

- [295] H. Gallaire, J. Minker, and J.-M. Nicolas (eds.). *Advances in Database Theory, Vols. 1 and 2*. Plenum Press, 1984.
- [296] H. Gallaire and J. Minker (eds.). *Logic and Data Bases*. Plenum Press, 1978.
- [297] S. Ganguly, W. Hasan, and R. Krishnamurthy. Query optimization for parallel execution. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [298] R. Ganski and H. Wong. Optimization of nested SQL queries revisited. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [299] V. Ganti, J. Gehrke, and R. Ramakrishnan. Demon: mining and monitoring evolving data. *IEEE Transactions on Knowledge and Data Engineering*, 13(1), 2001.
- [300] V. Ganti, J. Gehrke, R. Ramakrishnan, and W.-Y. Loh. Focus: a framework for measuring changes in data characteristics. In *Proc. ACM Symposium on Principles of Database Systems*, 1999.
- [301] V. Ganti, J. E. Gehrke, and R. Ramakrishnan. Cactus-clustering categorical data using summaries. In *Proc. ACM Intl. Conf. on Knowledge Discovery in Databases*, 1999.
- [302] V. Ganti, R. Ramakrishnan, J. E. Gehrke, A. Powell, and J. French. Clustering large datasets in arbitrary metric spaces. In *Proc. IEEE Intl. Conf. Data Engineering*, 1999.
- [303] H. Garcia-Molina and D. Barbara. How to assign votes in a distributed system. *Journal of the ACM*, 32(4), 1985.
- [304] H. Garcia-Molina, R. Lipton, and J. Valdes. A massive memory system machine. *IEEE Transactions on Computers*, C33(4):391–399, 1984.
- [305] H. Garcia-Molina, J. Ullman, and J. Widom. *Database Systems: The Complete Book* Prentice Hall, 2001.
- [306] H. Garcia-Molina and G. Wiederhold. Read-only transactions in a distributed database. *ACM Transactions on Database Systems*, 7(2):209–234, 1982.
- [307] E. Garfield. Citation analysis as a tool in journal evaluation. *Science*, 178(4060):471–479, 1972.
- [308] A. Garg and C. Gotlieb. Order preserving key transformations. *ACM Transactions on Database Systems*, 11(2):213–234, 1986.
- [309] J. E. Gehrke, V. Ganti, R. Ramakrishnan, and W.-Y. Loh. Boat: Optimistic decision tree construction. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1999.

- [310] J. E. Gehrke, F. Korn, and D. Srivastava. On computing correlated aggregates over continual data streams. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 2001.
- [311] J. E. Gehrke, R. Ramakrishnan, and V. Ganti. Rainforest: A framework for fast decision tree construction of large datasets. In *Proc. Intl. Conf. on Very Large Databases*, 1998.
- [312] S. P. Ghosh. *Data Base Organization for Data Management (2nd ed.)*. Academic Press, 1986.
- [313] P. B. Gibbons, Y. Matias, and V. Poosala. Fast incremental maintenance of approximate histograms. In *Proc. of the Conf. on Very Large Databases*, 1997.
- [314] P. B. Gibbons and Y. Matias. New sampling-based summary statistics for improving approximate query answers. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 331–342. ACM Press, 1998.
- [315] D. Gibson, J. M. Kleinberg, and P. Raghavan. Clustering categorical data: An approach based on dynamical systems. In *Proc. Intl. Conf. Very Large Data Bases*, 1998.
- [316] D. Gibson, J. M. Kleinberg, and P. Raghavan. Inferring web communities from link topology. In *Proc. ACM Conf. on Hypertext*, 1998.
- [317] G. A. Gibson. *Redundant Disk Arrays: Reliable, Parallel Secondary Storage*. An ACM Distinguished Dissertation 1991. MIT Press, 1992.
- [318] D. Gifford. Weighted voting for replicated data. In *ACM Symp. on Operating Systems Principles*, 1979.
- [319] A. C. Gilbert, Y. Kotidis, S. Muthukrishnan, and M. J. Strauss. Surfing wavelets on streams: One-pass summaries for approximate aggregate queries. In *Proc. of the Conf. on Very Large Databases*, 2001.
- [320] C. F. Goldfarb and P. Prescod. *The XML Handbook*. PrenticeHall, 1998.
- [321] R. Goldman and J. Widom. DataGuides: enabling query formulation and optimization in semistructured databases. In *Proc. Intl. Conf. on Very Large Data Bases*, pages 436–445, 1997.
- [322] J. Goldstein, R. Ramakrishnan, U. Shaft, and J.-B. Yu. Processing queries by linear constraints. In *Proc. ACM Symposium on Principles of Database Systems*, 1997.
- [323] G. Graefe. Encapsulation of parallelism in the Volcano query processing system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [324] G. Graefe. Query evaluation techniques for large databases. *ACM Computing Surveys*, 25(2), 1993.

- [325] G. Graefe, R. Bunker, and S. Cooper. Hash joins and hash teams in microsoft SQL Server: In *Proc. Intl. Conf. on Very Large Databases*, 1998.
- [326] G. Graefe and D. DeWitt. The Exodus optimizer generator. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [327] G. Graefe and K. Ward. Dynamic query optimization plans. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1989.
- [328] M. Graham, A. Mendelzon, and M. Vardi. Notions of dependency satisfaction. *Journal of the ACM*, 33(1):105–129, 1986.
- [329] G. Grahne. *The Problem of Incomplete Information in Relational Databases*. Springer-Verlag, 1991.
- [330] L. Gravano, H. Garcia-Molina, and A. Tomasic. Gloss: text-source discovery over the internet. *ACM Transactions on Database Systems*, 24(2), 1999.
- [331] J. Gray. Notes on data base operating systems. In *Operating Systems: An Advanced Course*. eds. Bayer, Graham, and Seegmuller, Springer-Verlag, 1978.
- [332] J. Gray. The transaction concept: Virtues and limitations. In *Proc. Intl. Conf. on Very Large Databases*, 1981.
- [333] J. Gray. Transparency in its place—the case against transparent access to geographically distributed data. *Tandem Computers, TR-89-1*, 1989.
- [334] J. Gray. *The Benchmark Handbook: for Database and Transaction Processing Systems*. Morgan Kaufmann, 1991.
- [335] J. Gray, A. Bosworth, A. Layman, and H. Pirahesh. Data cube: A relational aggregation operator generalizing group-by, cross-tab and sub-totals. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1996.
- [336] J. Gray, R. Lorie, G. Putzolu, and I. Traiger. Granularity of locks and degrees of consistency in a shared data base. In *Proc. of IFIP Working Conf. on Modelling of Data Base Management Systems*, 1977.
- [337] J. Gray, P. McJones, M. Blasgen, B. Lindsay, R. Lorie, G. Putzolu, T. Price, and I. Traiger. The recovery manager of the System R database manager. *ACM Computing Surveys*, 13(2):223–242, 1981.
- [338] J. Gray and A. Reuter. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann, 1992.
- [339] P. Gray. *Logic, Algebra, and Databases*. John Wiley, 1984.

- [340] M. Greenwald and S. Khanna. Space-efficient online computation of quantile summaries. In *Proc. ACM SIGMOD Conf. on Management of Data*, 2001.
- [341] P. Griffiths and B. Wade. An authorization mechanism for a relational database system. *ACM Transactions on Database Systems*, 1(3):242–255, 1976.
- [342] G. Grinstein. Visualization and data mining. In *Intl. Conf. on Knowledge Discovery in Databases*, 1996.
- [343] S. Guha, N. Mishra, R. Motwani, and L. O’Callaghan. Clustering data streams. In *Proc. of the Annual Symp. on Foundations of Computer Science*, 2000.
- [344] S. Guha, R. Rastogi, and K. Shim. Cure: an efficient clustering algorithm for large databases. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1998.
- [345] S. Guha, N. Koudas, and K. Shim. Data streams and histograms. In *Proc. of the ACM Symp. on Theory of Computing*, 2001.
- [346] D. Gunopulos, H. Mannila, R. Khardon, and H. Toivonen. Data mining, hypergraph transversals, and machine learning. In *Proc. ACM Symposium on Principles of Database Systems*, pages 209–216, 1997.
- [347] D. Gunopulos, H. Mannila, and S. Saluja. Discovering all most specific sentences by randomized algorithms. In *Proc. of the Intl. Conf. on Database Theory*, volume 1186 of *Lecture Notes in Computer Science*, pages 215–229. Springer, 1997.
- [348] A. Gupta and I. Mumick. *Materialized Views: Techniques, Implementations, and Applications* MIT Press, 1999.
- [349] A. Gupta, I. Mumick, and V. Subrahmanian. Maintaining views incrementally. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.
- [350] A. Guttman. R-trees: a dynamic index structure for spatial searching. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1984.
- [351] L. Haas, W. Chang, G. Lohman, J. McPherson, P. Wilms, G. Lapis, B. Lindsay, H. Pirahesh, M. Carey, and E. Shekita. Starburst mid-flight: As the dust clears. *IEEE Transactions on Knowledge and Data Engineering*, 2(1), 1990.
- [352] P. Haas, J. Naughton, S. Seshadri, and L. Stokes. Sampling-based estimation of the number of distinct values of an attribute. In *Proc. Intl. Conf. on Very Large Databases*, 1995.

- [353] P. Haas and A. Swami. Sampling-based selectivity estimation for joins using augmented frequent value statistics. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1995.
- [354] P. J. Haas and J. M. Hellerstein. Ripple joins for online aggregation. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 287–298. ACM Press, 1999.
- [355] T. Haerder and A. Reuter. Principles of transaction oriented database recovery—a taxonomy. *ACM Computing Surveys*, 15(4), 1982.
- [356] U. Halici and A. Dogac. Concurrency control in distributed databases through time intervals and short-term locks. *IEEE Transactions on Software Engineering*, 15(8):994–1003, 1989.
- [357] M. Hall. *Core Web Programming: HTML , Java , CGI , & Javascript*. Prentice-Hall, 1997.
- [358] P. Hall. Optimization of a simple expression in a relational data base system. *IBM Journal of Research and Development*, 20(3):244–257, 1976.
- [359] G. Hamilton, R. G. Cattell, and M. Fisher. *JDBC Database Access With Java: A Tutorial and Annotated Reference*. Java Series. Addison-Wesley, 1997.
- [360] M. Hammer and D. McLeod. Semantic integrity in a relational data base system. In *Proc. Intl. Conf. on Very Large Databases*, 1975.
- [361] J. Han and Y. Fu. Discovery of multiple-level association rules from large databases. In *Proc. Intl. Conf. on Very Large Databases*, 1995.
- [362] D. Hand. *Construction and Assessment of Classification Rules*. John Wiley & Sons, Chichester, England, 1997.
- [363] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, 2000.
- [364] J. Han, J. Pei, and Y. Yin. Mining frequent patterns without candidate generation. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, pages 1–12, 2000.
- [365] E. Hanson. A performance analysis of view materialization strategies. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [366] E. Hanson. Rule condition testing and action execution in Ariel. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [367] V. Harinarayan, A. Rajaraman, and J. Ullman. Implementing data cubes efficiently. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.

- [368] J. Haritsa, M. Carey, and M. Livny. On being optimistic about real-time constraints. In *ACM Symp. on Principles of Database Systems*, 1990.
- [369] J. Harrison and S. Dietrich. Maintenance of materialized views in deductive databases: An update propagation approach. In *Proc. Workshop on Deductive Databases*, 1992.
- [370] T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Verlag, 2001.
- [371] D. Heckerman. Bayesian networks for knowledge discovery. In *Advances in Knowledge Discovery and Data Mining*. eds. U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, MIT Press, 1996.
- [372] D. Heckerman, H. Mannila, D. Pregibon, and R. Uthurusamy, editors. *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1997.
- [373] J. Hellerstein. Optimization and execution techniques for queries with expensive methods. *Ph.D. thesis, University of Wisconsin-Madison*, 1995.
- [374] J. Hellerstein, P. Haas, and H. Wang. Online aggregation In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [375] J. Hellerstein, E. Koutsoupias, and C. Papadimitriou. On the analysis of indexing schemes. In *Proceedings of the ACM Symposium on Principles of Database Systems*, pages 249–256. ACM Press, 1997.
- [376] J. Hellerstein, J. Naughton, and A. Pfeffer. Generalized search trees for database systems. In *Proc. Intl. Conf. on Very Large Databases*, 1995.
- [377] J. M. Hellerstein, E. Koutsoupias, and C. H. Papadimitriou. On the analysis of indexing schemes. In *Proc. ACM Symposium on Principles of Database Systems*, pages 249–256, 1997.
- [378] C. Hidber. Online association rule mining. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 145–156, 1999.
- [379] R. Himmeroeder, G. Lausen, B. Ludaescher, and C. Schleppehorst. On a declarative semantics for Web queries. *Lecture Notes in Computer Science*, 1341:386–398, 1997.
- [380] C.-T. Ho, R. Agrawal, N. Megiddo, and R. Srikant. Range queries in OLAP data cubes. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [381] S. Holzner. *XML Complete*. Mc Graw-Hill, 1998.
- [382] D. Hong, T. Johnson, and U. Chakravarthy. Real-time transaction scheduling: A cost conscious approach. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.

- [383] W. Hong and M. Stonebraker. Optimization of parallel query execution plans in XPRS. In *Proc. Intl. Conf. on Parallel and Distributed Information Systems*, 1991.
- [384] W.-C. Hou and G. Ozsoyoglu. Statistical estimators for aggregate relational algebra queries. *ACM Transactions on Database Systems*, 16(4), 1991.
- [385] H. Hsiao and D. DeWitt. A performance study of three high availability data replication strategies. In *Proc. Intl. Conf. on Parallel and Distributed Information Systems*, 1991.
- [386] J. Huang, J. Stankovic, K. Ramamritham, and D. Towsley. Experimental evaluation of real-time optimistic concurrency control schemes. In *Proc. Intl. Conf. on Very Large Databases*, 1991.
- [387] Y. Huang, A. Sistla, and O. Wolfson. Data replication for mobile computers. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.
- [388] Y. Huang and O. Wolfson. A competitive dynamic data replication algorithm. In *Proc. IEEE CS IEEE Intl. Conf. on Data Engineering*, 1993.
- [389] R. Hull. Managing semantic heterogeneity in databases: A theoretical perspective. In *ACM Symp. on Principles of Database Systems*, 1997.
- [390] R. Hull and R. King. Semantic database modeling: Survey, applications, and research issues. *ACM Computing Surveys*, 19(19):201–260, 1987.
- [391] R. Hull and J. Su. Algebraic and calculus query languages for recursively typed complex objects. *Journal of Computer and System Sciences*, 47(1):121–156, 1993.
- [392] R. Hull and M. Yoshikawa. ILOG: Declarative creation and manipulation of object-identifiers. In *Proc. Intl. Conf. on Very Large Databases*, 1990.
- [393] G. Hulten, L. Spencer, and P. Domingos. Mining time-changing data streams. In *Proc. ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*, pages 97–106. AAAI Press, 2001.
- [394] J. Hunter. *Java Servlet Programming*. O'Reilly Associates, Inc., 1998.
- [395] T. Imielinski and H. Korth (eds.). *Mobile Computing*. Kluwer Academic, 1996.
- [396] T. Imielinski and W. Lipski. Incomplete information in relational databases. *Journal of the ACM*, 31(4):761–791, 1984.
- [397] T. Imielinski and H. Mannila. A database perspective on knowledge discovery. *Communications of the ACM*, 38(11):58–64, 1996.

- [398] T. Imielinski, S. Viswanathan, and B. Badrinath. Energy efficient indexing on air. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.
- [399] Y. Ioannidis. Query optimization. In *Handbook of Computer Science*. ed. A.B. Tucker, CRC Press, 1996.
- [400] Y. Ioannidis and S. Christodoulakis. Optimal histograms for limiting worst-case error propagation in the size of join results. *ACM Transactions on Database Systems*, 1993.
- [401] Y. Ioannidis and Y. Kang. Randomized algorithms for optimizing large join queries. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [402] Y. Ioannidis and Y. Kang. Left-deep vs. bushy trees: An analysis of strategy spaces and its implications for query optimization. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1991.
- [403] Y. Ioannidis, R. Ng, K. Shim, and T. Sellis. Parametric query processing. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [404] Y. Ioannidis and R. Ramakrishnan. Containment of conjunctive queries: Beyond relations as sets. *ACM Transactions on Database Systems*, 20(3):288–324, 1995.
- [405] Y. E. Ioannidis. Universality of serial histograms. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [406] H. Jagadish, D. Lieuwen, R. Rastogi, A. Silberschatz, and S. Sudarshan. Dali: A high performance main-memory storage manager. In *Proc. Intl. Conf. on Very Large Databases*, 1994.
- [407] A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. PrenticeHall, 1988.
- [408] S. Jajodia and D. Mutchler. Dynamic voting algorithms for maintaining the consistency of a replicated database. *ACM Transactions on Database Systems*, 15(2):230–280, 1990.
- [409] S. Jajodia and R. Sandhu. Polyinstantiation integrity in multilevel relations. In *Proc. IEEE Symp. on Security and Privacy*, 1990.
- [410] M. Jarke and J. Koch. Query optimization in database systems. *ACM Computing Surveys*, 16(2):111–152, 1984.
- [411] K. S. Jones and P. Willett, editors. *Readings in Information Retrieval. Multimedia Information and Systems*. Morgan Kaufmann Publishers, 1997.
- [412] J. Jou and P. Fischer. The complexity of recognizing 3NF schemes. *Information Processing Letters*, 14(4):187–190, 1983.

- [413] N. Kabra and D. J. DeWitt. Efficient mid-query re-optimization of sub-optimal query execution plans. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 1998.
- [414] Y. Kambayashi, M. Yoshikawa, and S. Yajima. Query processing for distributed databases using generalized semi-joins. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1982.
- [415] P. Kanellakis. Elements of relational database theory. In *Handbook of Theoretical Computer Science*. ed. J. Van Leeuwen, Elsevier, 1991.
- [416] P. Kanellakis. Constraint programming and database languages: A tutorial. In *ACM Symp. on Principles of Database Systems*, 1995.
- [417] H. Kargupta and P. Chan, editors. *Advances in Distributed and Parallel Knowledge Discovery*. MIT Press, 2000.
- [418] L. Kaufman and P. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley and Sons, 1990.
- [419] R. Kaushik, P. Bohannon, J. F. Naughton, and H. F. Korth. Covering indexes for branching path expression queries. In *Proceedings of SIGMOD*, 2002.
- [420] D. Keim and H.-P. Kriegel. VisDB: a system for visualizing large databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1995.
- [421] D. Keim and H.-P. Kriegel. Visualization techniques for mining large databases: A comparison. *IEEE Transactions on Knowledge and Data Engineering*, 8(6):923–938, 1996.
- [422] A. Keller. Algorithms for translating view updates to database updates for views involving selections, projections, and joins. *ACM Symp. on Principles of Database Systems*, 1985.
- [423] W. Kent. *Data and Reality, Basic Assumptions in Data Processing Reconsidered*. North-Holland, 1978.
- [424] W. Kent, R. Ahmed, J. Albert, M. Ketabchi, and M.-C. Shan. Object identification in multi-database systems. In *IFIP Intl. Conf. on Data Semantics*, 1992.
- [425] L. Kerschberg, A. Klug, and D. Tsichritzis. A taxonomy of data models. In *Systems for Large Data Bases*. eds. P.C. Lockemann and E.J. Neuhold, North-Holland, 1977.
- [426] W. Kiessling. On semantic reefs and efficient processing of correlation queries with aggregates. In *Proc. Intl. Conf. on Very Large Databases*, 1985.

- [427] M. Kifer, W. Kim, and Y. Sagiv. Querying object-oriented databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [428] M. Kifer, G. Lausen, and J. Wu. Logical foundations of object-oriented and frame-based languages. *Journal of the ACM*, 42(4):741–843, 1995.
- [429] M. Kifer and E. Lozinskii. Sygraf: Implementing logic programs in a database style. *IEEE Transactions on Software Engineering*, 14(7):922–935, 1988.
- [430] W. Kim. On optimizing an SQL -like nested query. *ACM Transactions on Database Systems*, 7(3), 1982.
- [431] W. Kim. Object-oriented database systems: Promise, reality, and future. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [432] W. Kim, J. Garza, N. Ballou, and D. Woelk. Architecture of the ORION next-generation database system. *IEEE Transactions on Knowledge and Data Engineering*, 2(1):109–124, 1990.
- [433] W. Kim and F. Lochovsky (eds.). *Object-Oriented Concepts, Databases, and Applications*. Addison-Wesley, 1989.
- [434] W. Kim, D. Reiner, and D. Batory (eds.). *Query Processing in Database Systems*. Springer Verlag, 1984.
- [435] W. Kim (ed.). *Modern Database Systems*. ACM Press and Addison-Wesley, 1995.
- [436] R. Kimball. *The Data Warehouse Toolkit*. John Wiley and Sons, 1996.
- [437] J. King. Quist: A system for semantic query optimization in relational databases. In *Proc. Intl. Conf. on Very Large Databases*, 1981.
- [438] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. In *Proc. ACM -SIAM Symp. on Discrete Algorithms*, 1998.
- [439] A. Klug. Equivalence of relational algebra and relational calculus query languages having aggregate functions. *Journal of the ACM*, 29(3):699–717, 1982.
- [440] A. Klug. On conjunctive queries containing inequalities. *Journal of the ACM*, 35(1):146–160, 1988.
- [441] E. Knapp. Deadlock detection in distributed databases. *ACM Computing Surveys*, 19(4):303–328, 1987.
- [442] D. Knuth. *The Art of Computer Programming, Vol.3—Sorting and Searching*. Addison-Wesley, 1973.
- [443] G. Koch and K. Loney. *Oracle: The Complete Reference*. Oracle Press, Osborne-McGraw-Hill, 1995.

- [444] W. Kohler. A survey of techniques for synchronization and recovery in decentralized computer systems. *ACM Computing Surveys*, 13(2):149–184, 1981.
- [445] D. Konopnicki and O. Shmueli. W3QS: A system for WWW querying. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1997.
- [446] F. Korn, H. Jagadish, and C. Faloutsos. Efficiently supporting ad hoc queries in large datasets of time sequences. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1997.
- [447] M. Kornacker, C. Mohan, and J. Hellerstein. Concurrency and recovery in generalized search trees. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [448] H. Korth, N. Soparkar, and A. Silberschatz. Triggered real-time databases with consistency constraints. In *Proc. Intl. Conf. on Very Large Databases*, 1990.
- [449] H. F. Korth. Deadlock freedom using edge locks. *ACM Transactions on Database Systems*, 7(4):632–652, 1982.
- [450] D. Kossmann. The state of the art in distributed query processing. *ACM Computing Surveys*, 32(4):422–469, 2000.
- [451] Y. Kotidis and N. Roussopoulos. An alternative storage organization for ROLAP aggregate views based on cubetrees. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 1998.
- [452] N. Krishnakumar and A. Bernstein. High throughput escrow algorithms for replicated databases. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [453] R. Krishnamurthy, H. Boral, and C. Zaniolo. Optimization of nonrecursive queries. In *Proc. Intl. Conf. on Very Large Databases*, 1986.
- [454] J. Kuhns. Logical aspects of question answering by computer. Technical report, Rand Corporation, RM-5428-Pr., 1967.
- [455] V. Kumar. *Performance of Concurrency Control Mechanisms in Centralized Database Systems*. PrenticeHall, 1996.
- [456] H. Kung and P. Lehman. Concurrent manipulation of binary search trees. *ACM Transactions on Database Systems*, 5(3):354–382, 1980.
- [457] H. Kung and J. Robinson. On optimistic methods for concurrency control. *Proc. Intl. Conf. on Very Large Databases*, 1979.
- [458] D. Kuo. Model and verification of a data manager based on ARIES. In *Intl. Conf. on Database Theory*, 1992.

- [459] M. LaCroix and A. Pirotte. Domain oriented relational languages. In *Proc. Intl. Conf. on Very Large Databases*, 1977.
- [460] M.-Y. Lai and W. Wilkinson. Distributed transaction management in Jasmin. In *Proc. Intl. Conf. on Very Large Databases*, 1984.
- [461] L. Lakshmanan, F. Sadri, and I. N. Subramanian. A declarative query language for querying and restructuring the web. In *Proc. Intl. Conf. on Research Issues in Data Engineering*, 1996.
- [462] L. V. S. Lakshmanan, Raymond T. Ng, J. Han, and A. Pang. Optimization of constrained frequent set queries with 2-variable constraints. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, pages 157–168. ACM Press, 1999.
- [463] C. Lam, G. Landis, J. Orenstein, and D. Weinreb. The Objectstore database system. *Communications of the ACM*, 34(10), 1991.
- [464] L. Lamport. Time, clocks and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.
- [465] B. Lampson and D. Lomet. A new presumed commit optimization for two phase commit. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [466] B. Lampson and H. Sturgis. Crash recovery in a distributed data storage system. Technical report, Xerox PARC, 1976.
- [467] C. Landwehr. Formal models of computer security. *ACM Computing Surveys*, 13(3):247–278, 1981.
- [468] R. Langerak. View updates in relational databases with an independent scheme. *ACM Transactions on Database Systems*, 15(1):40–66, 1990.
- [469] P.-A. Larson. Linear hashing with overflow-handling by linear probing. *ACM Transactions on Database Systems*, 10(1):75–89, 1985.
- [470] P.-A. Larson. Linear hashing with separators—A dynamic hashing scheme achieving one-access retrieval. *ACM Transactions on Database Systems*, 13(3):366–388, 1988.
- [471] P.-A. Larson and G. Graefe. Memory Management During Run Generation in External Sorting. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1998.
- [472] P. Lehman and S. Yao. Efficient locking for concurrent operations on b trees. *ACM Transactions on Database Systems*, 6(4):650–670, 1981.
- [473] T. Leung and R. Muntz. Temporal query processing and optimization in multiprocessor database machines. In *Proc. Intl. Conf. on Very Large Databases*, 1992.

- [474] M. Leventhal, D. Lewis, and M. Fuchs. *Designing XML Internet applications*. The Charles F. Goldfarb series on open information management. PrenticeHall, 1998.
- [475] P. Lewis, A. Bernstein, and M. Kifer. *Databases and Transaction Processing*. Addison Wesley, 2001.
- [476] E.-P. Lim and J. Srivastava. Query optimization and processing in federated database systems. In *Proc. Intl. Conf. on Intelligent Knowledge Management*, 1993.
- [477] B. Lindsay, J. McPherson, and H. Pirahesh. A data management extension architecture. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [478] B. Lindsay, P. Selinger, C. Galtieri, J. Gray, R. Lorie, G. Putzolu, I. Traiger, and B. Wade. Notes on distributed databases. Technical report, RJ2571, San Jose, CA, 1979.
- [479] D.-I. Lin and Z. M. Kedem. Pincer search: A new algorithm for discovering the maximum frequent set. *Lecture Notes in Computer Science*, 1377:105–??, 1998.
- [480] V. Linnemann, K. Kuspert, P. Dadam, P. Pistor, R. Erbe, A. Kemper, N. Sudkamp, G. Walch, and M. Wallrath. Design and implementation of an extensible database management system supporting user defined data types and functions. In *Proc. Intl. Conf. on Very Large Databases*, 1988.
- [481] R. Lipton, J. Naughton, and D. Schneider. Practical selectivity estimation through adaptive sampling. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [482] B. Liskov, A. Adya, M. Castro, M. Day, S. Ghemawat, R. Gruber, U. Maheshwari, A. Myers, and L. Shrira. Safe and efficient sharing of persistent objects in Thor. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [483] W. Litwin. Linear Hashing: A new tool for file and table addressing. In *Proc. Intl. Conf. on Very Large Databases*, 1980.
- [484] W. Litwin. Trie Hashing. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1981.
- [485] W. Litwin and A. Abdellatif. Multidatabase interoperability. *IEEE Computer*, 12(19):10–18, 1986.
- [486] W. Litwin, L. Mark, and N. Roussopoulos. Interoperability of multiple autonomous databases. *ACM Computing Surveys*, 22(3), 1990.

- [487] W. Litwin, M.-A. Neimat, and D. Schneider. LH \*—A scalable, distributed data structure. *ACM Transactions on Database Systems*, 21(4):480–525, 1996.
- [488] M. Liu, A. Sheth, and A. Singhal. An adaptive concurrency control strategy for distributed database system. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1984.
- [489] M. Livny, R. Ramakrishnan, K. Beyer, G. Chen, D. Donjerkovic, S. Lawande, J. Myllymaki, and K. Wenger. DEVise: Integrated querying and visual exploration of large datasets. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [490] G. Lohman. Grammar-like functional rules for representing query optimization alternatives. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1988.
- [491] D. Lomet and B. Salzberg. The hB-T tree: A multiattribute indexing method with good guaranteed performance. *ACM Transactions on Database Systems*, 15(4), 1990.
- [492] D. Lomet and B. Salzberg. Access method concurrency with recovery. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [493] R. Lorie. Physical integrity in a large segmented database. *ACM Transactions on Database Systems*, 2(1):91–104, 1977.
- [494] R. Lorie and H. Young. A low communication sort algorithm for a parallel database machine. In *Proc. Intl. Conf. on Very Large Databases*, 1989.
- [495] Y. Lou and Z. Ozsoyoglu. LLO: An object-oriented deductive language with methods and method inheritance. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1991.
- [496] H. Lu, B.-C. Ooi, and K.-L. Tan (eds.). *Query Processing in Parallel Relational Database Systems*. IEEE Computer Society Press, 1994.
- [497] C. Lucchesi and S. Osborn. Candidate keys for relations. *J. Computer and System Sciences*, 17(2):270–279, 1978.
- [498] V. Lum. Multi-attribute retrieval with combined indexes. *Communications of the ACM*, 1(11):660–665, 1970.
- [499] T. Lunt, D. Denning, R. Schell, M. Heckman, and W. Shockley. The seaview security model. *IEEE Transactions on Software Engineering*, 16(6):593–607, 1990.
- [500] L. Mackert and G. Lohman. R\* optimizer validation and performance evaluation for local queries. Technical report, IBM RJ-4989, San Jose, CA, 1986.

- [501] D. Maier. *The Theory of Relational Databases*. Computer Science Press, 1983.
- [502] D. Maier, A. Mendelzon, and Y. Sagiv. Testing implication of data dependencies. *ACM Transactions on Database Systems*, 4(4), 1979.
- [503] D. Maier and D. Warren. *Computing with Logic: Logic Programming with Prolog*. Benjamin/Cummings Publishers, 1988.
- [504] A. Makinouchi. A consideration on normal form of not-necessarily-normalized relation in the relational data model. In *Proc. Intl. Conf. on Very Large Databases*, 1977.
- [505] U. Manber and R. Ladner. Concurrency control in a dynamic search structure. *ACM Transactions on Database Systems*, 9(3):439–455, 1984.
- [506] G. Manku, S. Rajagopalan, and B. Lindsay. Random sampling techniques for space efficient online computation of order statistics of large datasets. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1999.
- [507] H. Mannila. Methods and problems in data mining. In *Intl. Conf. on Database Theory*, 1997.
- [508] H. Mannila and K.-J. Raiha. Design by Example: An application of Armstrong relations. *Journal of Computer and System Sciences*, 33(2):126–141, 1986.
- [509] H. Mannila and K.-J. Raiha. *The Design of Relational Databases*. Addison-Wesley, 1992.
- [510] H. Mannila, H. Toivonen, and A. I. Verkamo. Discovering frequent episodes in sequences. In *Proc. Intl. Conf. on Knowledge Discovery in Databases and Data Mining*, 1995.
- [511] H. Mannila, P. Smyth, and D. J. Hand. *Principles of Data Mining*. MIT Press, 2001.
- [512] M. Mannino, P. Chu, and T. Sager. Statistical profile estimation in database systems. *ACM Computing Surveys*, 20(3):191–221, 1988.
- [513] V. Markowitz. Representing processes in the extended entity-relationship model. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1990.
- [514] V. Markowitz. Safe referential integrity structures in relational databases. In *Proc. Intl. Conf. on Very Large Databases*, 1991.
- [515] Y. Matias, J. S. Vitter, and M. Wang. Dynamic maintenance of wavelet-based histograms. In *Proc. of the Conf. on Very Large Databases*, 2000.
- [516] D. McCarthy and U. Dayal. The architecture of an active data base management system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1989.

- [517] W. McCune and L. Henschen. Maintaining state constraints in relational databases: A proof theoretic basis. *Journal of the ACM*, 36(1):46–68, 1989.
- [518] J. McHugh, S. Abiteboul, R. Goldman, D. Quass, and J. Widom. Lore: A database management system for semistructured data. *ACM SIGMOD Record*, 26(3):54–66, 1997.
- [519] S. Mehrotra, R. Rastogi, Y. Breitbart, H. Korth, and A. Silberschatz. Ensuring transaction atomicity in multidatabase systems. In *ACM Symp. on Principles of Database Systems*, 1992.
- [520] S. Mehrotra, R. Rastogi, H. Korth, and A. Silberschatz. The concurrency control problem in multidatabases: Characteristics and solutions. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [521] M. Mehta, R. Agrawal, and J. Rissanen. SLIQ: A fast scalable classifier for data mining. In *Proc. Intl. Conf. on Extending Database Technology*, 1996.
- [522] M. Mehta, V. Soloviev, and D. DeWitt. Batch scheduling in parallel database systems. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1993.
- [523] J. Melton. *Advanced SQL:1999, Understanding Understanding Object-Relational and Other Advanced Features*. Morgan Kaufmann, 2002.
- [524] J. Melton and A. Simon. *Understanding the New SQL: A Complete Guide*. Morgan Kaufmann, 1993.
- [525] J. Melton and A. Simon. *SQL:1999, Understanding Relational Language Components*. Morgan Kaufmann, 2002.
- [526] D. Menasce and R. Muntz. Locking and deadlock detection in distributed data bases. *IEEE Transactions on Software Engineering*, 5(3):195–222, 1979.
- [527] A. Mendelzon and T. Milo. Formal models of web queries. In *ACM Symp. on Principles of Database Systems*, 1997.
- [528] A. O. Mendelzon, G. A. Mihaila, and T. Milo. Querying the World Wide Web. *Journal on Digital Libraries*, 1:54–67, 1997.
- [529] R. Meo, G. Psaila, and S. Ceri. A new SQL -like operator for mining association rules. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [530] T. Merrett. The extended relational algebra, a basis for query languages. In *Databases*. ed. Shneiderman, Academic Press, 1978.
- [531] T. Merrett. *Relational Information Systems*. Reston Publishing Company, 1983.

- [532] D. Michie, D. Spiegelhalter, and C. Taylor, editors. *Machine Learning, Neural and Statistical Classification*. Ellis Horwood, London, 1994.
- [533] Microsoft. *Microsoft ODBC 3.0 Software Development Kit and Programmer's Reference*. Microsoft Press, 1997.
- [534] K. Mikkilineni and S. Su. An evaluation of relational join algorithms in a pipelined query processing environment. *IEEE Transactions on Software Engineering*, 14(6):838–848, 1988.
- [535] R. Miller, Y. Ioannidis, and R. Ramakrishnan. The use of information capacity in schema integration and translation. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [536] T. Milo and D. Suciu. Index structures for path expressions. In *ICDT: 7th International Conference on Database Theory*, 1999.
- [537] J. Minker (ed.). *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufmann, 1988.
- [538] T. Minoura and G. Wiederhold. Resilient extended true-copy token scheme for a distributed database. *IEEE Transactions in Software Engineering*, 8(3):173–189, 1982.
- [539] G. Mitchell, U. Dayal, and S. Zdonik. Control of an extensible query optimizer: A planning-based approach. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [540] A. Moffat and J. Zobel. Self-indexing inverted files for fast text retrieval. *ACM Transactions on Information Systems*, 14(4):349–379, 1996.
- [541] C. Mohan. ARIES/NT: A recovery method based on write-ahead logging for nested. In *Proc. Intl. Conf. on Very Large Databases*, 1989.
- [542] C. Mohan. Commit LSN: A novel and simple method for reducing locking and latching in transaction processing systems. In *Proc. Intl. Conf. on Very Large Databases*, 1990.
- [543] C. Mohan. ARIES/LHS: A concurrency control and recovery method using write-ahead logging for linear hashing with separators. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1993.
- [544] C. Mohan, D. Haderle, B. Lindsay, H. Pirahesh, and P. Schwarz. ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging. *ACM Transactions on Database Systems*, 17(1):94–162, 1992.
- [545] C. Mohan and F. Levine. ARIES/IM An efficient and high concurrency index management method using write-ahead logging. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.

- [546] C. Mohan and B. Lindsay. Efficient commit protocols for the tree of processes model of distributed transactions. In *ACM SIGACT-SIGOPS Symp. on Principles of Distributed Computing*, 1983.
- [547] C. Mohan, B. Lindsay, and R. Obermarck. Transaction management in the R\* distributed database management system. *ACM Transactions on Database Systems*, 11(4):378–396, 1986.
- [548] C. Mohan and I. Narang. Algorithms for creating indexes for very large tables without quiescing updates. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [549] K. Morris, J. Naughton, Y. Saraiya, J. Ullman, and A. Van Gelder. YAWN ! (Yet Another Window on NAIL! ). *Database Engineering*, 6:211–226, 1987.
- [550] A. Motro. Superviews: Virtual integration of multiple databases. *IEEE Transactions on Software Engineering*, 13(7):785–798, 1987.
- [551] A. Motro and P. Buneman. Constructing superviews. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1981.
- [552] R. Mukkamala. Measuring the effect of data distribution and replication models on performance evaluation of distributed database systems. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1989.
- [553] I. Mumick, S. Finkelstein, H. Pirahesh, and R. Ramakrishnan. Magic is relevant. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [554] I. Mumick, S. Finkelstein, H. Pirahesh, and R. Ramakrishnan. Magic conditions. *ACM Transactions on Database Systems*, 21(1):107–155, 1996.
- [555] I. Mumick, H. Pirahesh, and R. Ramakrishnan. Duplicates and aggregates in deductive databases. In *Proc. Intl. Conf. on Very Large Databases*, 1990.
- [556] I. Mumick and K. Ross. Noodle: A language for declarative querying in an object-oriented database. In *Intl. Conf. on Deductive and Object-Oriented Databases*, 1993.
- [557] M. Muralikrishna. Improved unnesting algorithms for join aggregate SQL queries. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [558] M. Muralikrishna and D. DeWitt. Equi-depth histograms for estimating selectivity factors for multi-dimensional queries. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1988.
- [559] S. Naqvi. Negation as failure for first-order queries. In *ACM Symp. on Principles of Database Systems*, 1986.

- [560] M. Negri, G. Pelagatti, and L. Sbattella. Formal semantics of SQL queries. *ACM Transactions on Database Systems*, 16(3), 1991.
- [561] S. Nestorov, J. Ullman, J. Weiner, and S. Chawathe. Representative objects: Concise representations of semistructured, hierarchical data. In *Proc. Intl. Conf. on Data Engineering*. IEEE Computer Society, 1997.
- [562] R. T. Ng and J. Han. Efficient and effective clustering methods for spatial data mining. In *Proc. Intl. Conf. on Very Large Databases*, Santiago, Chile, September 1994.
- [563] R. T. Ng, L. V. S. Lakshmanan, J. Han, and A. Pang. Exploratory mining and pruning optimizations of constrained association rules. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, pages 13–24. ACM Press, 1998.
- [564] T. Nguyen and V. Srinivasan. Accessing relational databases from the World Wide Web. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [565] J. Nievergelt, H. Hinterberger, and K. Sevcik. The Grid File: An adaptable symmetric multikey file structure. *ACM Transactions on Database Systems*, 9(1):38–71, 1984.
- [566] C. Nyberg, T. Barclay, Z. Cvetanovic, J. Gray, and D. Lomet. Alphasort: a cache-sensitive parallel external sort. *VLDB Journal*, 4(4):603–627, 1995.
- [567] R. Obermarck. Global deadlock detection algorithm. *ACM Transactions on Database Systems*, 7(2):187–208, 1981.
- [568] L. O’Callaghan, N. Mishra, A. Meyerson, S. Guha, and R. Motwani. Streaming-data algorithms for high-quality clustering. In *Proc. of the Intl. Conference on Data Engineering*. IEEE, 2002.
- [569] F. Olken and D. Rotem. Simple random sampling from relational databases. In *Proc. Intl. Conf. on Very Large Databases*, 1986.
- [570] F. Olken and D. Rotem. Maintenance of materialized views of sampling queries. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1992.
- [571] C. Olston, B. T. Loo, and J. Widom. Adaptive precision setting for cached approximate values. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 2001.
- [572] C. Olston and J. Widom. Offering a precision-performance tradeoff for aggregation queries over replicated data. In *Proc. of the Conf. on Very Large Databases*, pages 144–155, 2000.
- [573] C. Olston and J. Widom. Best-effort cache synchronization with source cooperation. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 2002.

- [574] P. O’Neil and E. O’Neil. *Database Principles, Programming, and Performance*. Addison Wesley, 2 edition, 2000.
- [575] P. O’Neil and D. Quass. Improved query performance with variant indexes. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [576] B. Ozden, R. Rastogi, and A. Silberschatz. Multimedia support for databases. In *ACM Symp. on Principles of Database Systems*, 1997.
- [577] G. Ozsoyoglu, K. Du, S. Guruswamy, and W.-C. Hou. Processing real-time, non-aggregate queries with time-constraints in case-db. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1992.
- [578] G. Ozsoyoglu, Z. Ozsoyoglu, and V. Matos. Extending relational algebra and relational calculus with set-valued attributes and aggregate functions. *ACM Transactions on Database Systems*, 12(4):566–592, 1987.
- [579] Z. Ozsoyoglu and L.-Y. Yuan. A new normal form for nested relations. *ACM Transactions on Database Systems*, 12(1):111–136, 1987.
- [580] M. Ozsu and P. Valduriez. *Principles of Distributed Database Systems*. PrenticeHall, 1991.
- [581] C. Papadimitriou. The serializability of concurrent database updates. *Journal of the ACM*, 26(4):631–653, 1979.
- [582] C. Papadimitriou. *The Theory of Database Concurrency Control*. Computer Science Press, 1986.
- [583] Y. Papakonstantinou, S. Abiteboul, and H. Garcia-Molina. Object fusion in mediator systems. In *Proc. Intl. Conf. on Very Large Data Bases*, 1996.
- [584] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object exchange across heterogeneous information sources. In *Proc. Intl. Conf. on Data Engineering*, 1995.
- [585] J. Park and A. Segev. Using common subexpressions to optimize multiple queries. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1988.
- [586] J. Patel, J.-B. Yu, K. Tufte, B. Nag, J. Burger, N. Hall, K. Ramasamy, R. Lueder, C. Ellman, J. Kupsch, S. Guo, D. DeWitt, and J. Naughton. Building a scaleable geo-spatial DBMS: Technology, implementation, and evaluation. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [587] D. Patterson, G. Gibson, and R. Katz. RAID: redundant arrays of inexpensive disks. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1988.

- [588] H.-B. Paul, H.-J. Schek, M. Scholl, G. Weikum, and U. Deppisch. Architecture and implementation of the Darmstadt database kernel system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [589] J. Peckham and F. Maryanski. Semantic data models. *ACM Computing Surveys*, 20(3):153–189, 1988.
- [590] J. Pei and J. Han. Can we push more constraints into frequent pattern mining? In *ACM SIGKDD Conference*, pages 350–354, 2000.
- [591] J. Pei, J. Han, and L. V. S. Lakshmanan. Mining frequent item sets with convertible constraints. In *Proc. Intl. Conf. on Data Engineering (ICDE)*, pages 433–442. IEEE Computer Society, 2001.
- [592] E. Petajan, Y. Jean, D. Lieuwen, and V. Anupam. DataSpace: An automated visualization system for large databases. In *Proc. of SPIE, Visual Data Exploration and Analysis*, 1997.
- [593] S. Petrov. Finite axiomatization of languages for representation of system properties. *Information Sciences*, 47:339–372, 1989.
- [594] G. Piatetsky-Shapiro and C. Cornell. Accurate estimation of the number of tuples satisfying a condition. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1984.
- [595] G. Piatetsky-Shapiro and W. J. Frawley, editors. *Knowledge Discovery in Databases*. AAAI/MIT Press, Menlo Park, CA, 1991.
- [596] H. Pirahesh and J. Hellerstein. Extensible/rule-based query rewrite optimization in starburst. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [597] N. Pitts-Moultis and C. Kirk. *XML black book: Indispensable problem solver*. Coriolis Group, 1998.
- [598] V. Poosala, Y. Ioannidis, P. Haas, and E. Shekita. Improved histograms for selectivity estimation of range predicates. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [599] C. Pu. Superdatabases for composition of heterogeneous databases. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1988.
- [600] C. Pu and A. Leff. Replica control in distributed systems: An asynchronous approach. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1991.
- [601] X.-L. Qian and G. Wiederhold. Incremental recomputation of active relational expressions. *IEEE Transactions on Knowledge and Data Engineering*, 3(3):337–341, 1990.

- [602] D. Quass, A. Rajaraman, Y. Sagiv, and J. Ullman. Querying semistructured heterogeneous information. In *Proc. Intl. Conf. on Deductive and Object-Oriented Databases*, 1995.
- [603] J. R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufman, 1993.
- [604] H. G. M. R. Alonso, D. Barbara. Data caching issues in an information retrieval system. *ACM Transactions on Database Systems*, 15(3), 1990.
- [605] The RAIDBook: A source book for RAID technology. The RAID Advisory Board, <http://www.raid-advisory.com>, North Grafton, MA, Dec. 1998. Sixth Edition.
- [606] D. Rafiei and A. Mendelzon. Similarity-based queries for time series data. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [607] M. Ramakrishna. An exact probability model for finite hash tables. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1988.
- [608] M. Ramakrishna and P.-A. Larson. File organization using composite perfect hashing. *ACM Transactions on Database Systems*, 14(2):231–263, 1989.
- [609] I. Ramakrishnan, P. Rao, K. Sagonas, T. Swift, and D. Warren. Efficient tabling mechanisms for logic programs. In *Intl. Conf. on Logic Programming*, 1995.
- [610] R. Ramakrishnan, D. Donjerkovic, A. Ranganathan, K. Beyer, and M. Krishnaprasad. SRQL: Sorted relational query language In *Proc. IEEE Intl. Conf. on Scientific and Statistical DBMS*, 1998.
- [611] R. Ramakrishnan, D. Srivastava, and S. Sudarshan. Efficient bottom-up evaluation of logic programs. In *The State of the Art in Computer Systems and Software Engineering*. ed. J. Vandewalle, Kluwer Academic, 1992.
- [612] R. Ramakrishnan, D. Srivastava, S. Sudarshan, and P. Seshadri. The CORAL: deductive system. *VLDB Journal*, 3(2):161–210, 1994.
- [613] R. Ramakrishnan, S. Stolfo, R. J. Bayardo., and I. Parsa, editors. *Proc. ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*. AAAI Press, 2000.
- [614] R. Ramakrishnan and J. Ullman. A survey of deductive database systems. *Journal of Logic Programming*, 23(2):125–149, 1995.
- [615] K. Ramamohanarao. Design overview of the Aditi deductive database system. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1991.

- [616] K. Ramamohanarao, J. Shepherd, and R. Sacks-Davis. Partial-match retrieval for dynamic files using linear hashing with partial expansions. In *Intl. Conf. on Foundations of Data Organization and Algorithms*, 1989.
- [617] V. Raman, B. Raman, and J. M. Hellerstein. Online dynamic reordering for interactive data processing. In *Proc. of the Conf. on Very Large Databases*, pages 709–720. Morgan Kaufmann, 1999.
- [618] S. Rao, A. Badia, and D. Van Gucht. Providing better support for a class of decision support queries. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [619] R. Rastogi and K. Shim. Public: A decision tree classifier that integrates building and pruning. In *Proc. Intl. Conf. on Very Large Databases*, 1998.
- [620] D. Reed. Implementing atomic actions on decentralized data. *ACM Transactions on Database Systems*, 1(1):3–23, 1983.
- [621] G. Reese. *Database Programming With JDBC and Java*. O’Reilly & Associates, 1997.
- [622] R. Reiter. A sound and sometimes complete query evaluation algorithm for relational databases with null values. *Journal of the ACM*, 33(2):349–370, 1986.
- [623] E. Rescorla. *SSL and TLS: Designing and Building Secure Systems*. Addison Wesley Professional, 2000.
- [624] A. Reuter. A fast transaction-oriented logging scheme for undo recovery. *IEEE Transactions on Software Engineering*, 6(4):348–356, 1980.
- [625] A. Reuter. Performance analysis of recovery techniques. *ACM Transactions on Database Systems*, 9(4):526–559, 1984.
- [626] E. Riloff and L. Hollaar. Text databases and information retrieval. In *Handbook of Computer Science*. ed. A.B. Tucker, CRC Press, 1996.
- [627] J. Rissanen. Independent components of relations. *ACM Transactions on Database Systems*, 2(4):317–325, 1977.
- [628] R. Rivest. Partial match retrieval algorithms. *SIAM Journal on Computing*, 5(1):19–50, 1976.
- [629] R. L. Rivest, A. Shamir, and L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.
- [630] J. T. Robinson. The KDB tree: A search structure for large multidimensional dynamic indexes. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, 1981.

- [631] J. Rohmer, F. Lescoeur, and J. Kerisit. The Alexander method, a technique for the processing of recursive queries. *New Generation Computing*, 4(3):273–285, 1986.
- [632] D. Rosenkrantz, R. Stearns, and P. Lewis. System level concurrency control for distributed database systems. *ACM Transactions on Database Systems*, 3(2), 1978.
- [633] A. Rosenthal and U. Chakravarthy. Anatomy of a modular multiple query optimizer. In *Proc. Intl. Conf. on Very Large Databases*, 1988.
- [634] K. Ross and D. Srivastava. Fast computation of sparse datacubes. In *Proc. Intl. Conf. on Very Large Databases*, 1997.
- [635] K. Ross, D. Srivastava, and S. Sudarshan. Materialized view maintenance and integrity constraint checking: Trading space for time. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [636] J. Rothnie, P. Bernstein, S. Fox, N. Goodman, M. Hammer, T. Landers, C. Reeve, D. Shipman, and E. Wong. Introduction to a system for distributed databases (SDD -1). *ACM Transactions on Database Systems*, 5(1), 1980.
- [637] J. Rothnie and N. Goodman. An overview of the preliminary design of SDD -1: A system for distributed data bases. In *Proc. Berkeley Workshop on Distributed Data Management and Computer Networks*, 1977.
- [638] N. Roussopoulos, Y. Kotidis, and M. Roussopoulos. Cubetree: Organization of and bulk updates on the data cube. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [639] S. Rozen and D. Shasha. Using feature set compromise to automate physical database design. In *Proc. Intl. Conf. on Very Large Databases*, 1991.
- [640] J. Rumbaugh, I. Jacobson, and G. Booch. *The Unified Modeling Language Reference Manual (Addison-Wesley Object Technology Series)*. Addison-Wesley, 1998.
- [641] M. Rusinkiewicz, A. Sheth, and G. Karabatis. Specifying interdatabase dependencies in a multidatabase environment. *IEEE Computer*, 24(12), 1991.
- [642] D. Sacca and C. Zaniolo. Magic counting methods. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [643] Y. Sagiv and M. Yannakakis. Equivalence among expressions with the union and difference operators. *Journal of the ACM*, 27(4):633–655, 1980.
- [644] K. Sagonas, T. Swift, and D. Warren. XSB as an efficient deductive database engine. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.

- [645] A. Sahuguet, L. Dupont, and T. Nguyen. Kweelt: Querying XML in the new millenium. <http://kweelt.sourceforge.net>, Sept 2000.
- [646] G. Salton and M. J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, 1983.
- [647] B. Salzberg, A. Tsukerman, J. Gray, M. Stewart, S. Uren, and B. Vaughan. Fastsort: A distributed single-input single-output external sort. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1990.
- [648] B. J. Salzberg. *File Structures*. PrenticeHall, 1988.
- [649] H. Samet. The Quad T tree and related hierarchical data structures. *ACM Computing Surveys*, 16(2), 1984.
- [650] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, 1990.
- [651] J. Sander, M. Ester, H.-P. Kriegel, and X. Xu. Density-based clustering in spatial databases. *J. of Data Mining and Knowledge Discovery*, 2(2), 1998.
- [652] R. E. Sanders. *ODBC 3.5 Developer's Guide*. McGraw-Hill Series on Data Warehousing and Data Management. McGraw-Hill, 1998.
- [653] S. Sarawagi and M. Stonebraker. Efficient organization of large multidimensional arrays. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1994.
- [654] S. Sarawagi, S. Thomas, and R. Agrawal. Integrating mining with relational database systems: Alternatives and implications. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 1998.
- [655] A. Savasere, E. Omiecinski, and S. Navathe. An efficient algorithm for mining association rules in large databases. In *Proc. Intl. Conf. on Very Large Databases*, 1995.
- [656] P. Schauble. Spider: A multiuser information retrieval system for semistructured and dynamic data. In *Proc. ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 318 – 327, 1993.
- [657] H.-J. Schek, H.-B. Paul, M. Scholl, and G. Weikum. The DASDBS project: Objects, experiences, and future projects. *IEEE Transactions on Knowledge and Data Engineering*, 2(1), 1990.
- [658] M. Schkolnick. Physical database design techniques. In *NYU Symp. on Database Design*, 1978.
- [659] M. Schkolnick and P. Sorenson. The effects of denormalization on database performance. Technical report, IBM RJ3082, San Jose, CA, 1981.

- [660] G. Schlageter. Optimistic methods for concurrency control in distributed database systems. In *Proc. Intl. Conf. on Very Large Databases*, 1981.
- [661] B. Schneier. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, 1995.
- [662] E. Sciore. A complete axiomatization of full join dependencies. *Journal of the ACM*, 29(2):373–393, 1982.
- [663] E. Sciore, M. Siegel, and A. Rosenthal. Using semantic values to facilitate interoperability among heterogeneous information systems. *ACM Transactions on Database Systems*, 19(2):254–290, 1994.
- [664] A. Segev and J. Park. Maintaining materialized views in distributed databases. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1989.
- [665] A. Segev and A. Shoshani. Logical modeling of temporal data. *Proc. ACM SIGMOD Conf. on the Management of Data*, 1987.
- [666] P. Selfridge, D. Srivastava, and L. Wilson. IDEA: Interactive data exploration and analysis. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [667] P. Selinger and M. Adiba. Access path selections in distributed database management systems. In *Proc. Intl. Conf. on Databases, British Computer Society*, 1980.
- [668] P. Selinger, M. Astrahan, D. Chamberlin, R. Lorie, and T. Price. Access path selection in a relational database management system. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1979.
- [669] T. K. Sellis. Multiple query optimization. *ACM Transactions on Database Systems*, 13(1):23–52, 1988.
- [670] P. Seshadri, J. Hellerstein, H. Pirahesh, T. Leung, R. Ramakrishnan, D. Srivastava, P. Stuckey, and S. Sudarshan. Cost-based optimization for Magic: Algebra and implementation. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1996.
- [671] P. Seshadri, M. Livny, and R. Ramakrishnan. The design and implementation of a sequence database system. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [672] P. Seshadri, M. Livny, and R. Ramakrishnan. The case for enhanced abstract data types. In *Proc. Intl. Conf. on Very Large Databases*, 1997.
- [673] P. Seshadri, H. Pirahesh, and T. Leung. Complex query decorrelation. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1996.
- [674] J. Shafer and R. Agrawal. SPRINT: a scalable parallel classifier for data mining. In *Proc. Intl. Conf. on Very Large Databases*, 1996.

- [675] J. Shanmugasundaram, U. Fayyad, and P. Bradley. Compressed data cubes for olap aggregate query approximation on continuous dimensions. In *Proc. Intl. Conf. on Knowledge Discovery and Data Mining (KDD)*, 1999.
- [676] J. Shanmugasundaram, J. Kiernan, E. J. Shekita, C. Fan, and J. Funderburk. Querying XML views of relational data. In *Proc. Intl. Conf. on Very Large Data Bases*, 2001.
- [677] L. Shapiro. Join processing in database systems with large main memories. *ACM Transactions on Database Systems*, 11(3):239–264, 1986.
- [678] D. Shasha and N. Goodman. Concurrent search structure algorithms. *ACM Transactions on Database Systems*, 13:53–90, 1988.
- [679] D. Shasha, E. Simon, and P. Valduriez. Simple rational guidance for chopping up transactions. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [680] H. Shatkay and S. Zdonik. Approximate queries and representations for large data sequences. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1996.
- [681] T. Sheard and D. Stemple. Automatic verification of database transaction safety. *ACM Transactions on Database Systems*, 1989.
- [682] S. Shenoy and Z. Ozsoyoglu. Design and implementation of a semantic query optimizer. *IEEE Transactions on Knowledge and Data Engineering*, 1(3):344–361, 1989.
- [683] P. Shenoy, J. Haritsa, S. Sudarshan, G. Bhalotia, M. Bawa, and D. Shah. Turbo-charging vertical mining of large databases. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, pages 22–33, May 2000.
- [684] A. Sheth and J. Larson. Federated database systems for managing distributed, heterogeneous, and autonomous databases. *Computing Surveys*, 22(3):183–236, 1990.
- [685] A. Sheth, J. Larson, A. Cornelio, and S. Navathe. A tool for integrating conceptual schemas and user views. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1988.
- [686] A. Shoshani. OLAP and statistical databases: Similarities and differences. In *ACM Symp. on Principles of Database Systems*, 1997.
- [687] A. Shukla, P. Deshpande, J. Naughton, and K. Ramasamy. Storage estimation for multidimensional aggregates in the presence of hierarchies. In *Proc. Intl. Conf. on Very Large Databases*, 1996.

- [688] M. Siegel, E. Sciore, and S. Salveter. A method for automatic rule derivation to support semantic query optimization. *ACM Transactions on Database Systems*, 17(4), 1992.
- [689] A. Silberschatz, H. Korth, and S. Sudarshan. *Database System Concepts (4th ed.)*. McGraw-Hill, 4 edition, 2001.
- [690] E. Simon, J. Kiernan, and C. de Maindreville. Implementing high-level active rules on top of relational databases. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [691] E. Simoudis, J. Wei, and U. M. Fayyad, editors. *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1996.
- [692] D. Skeen. Nonblocking commit protocols. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1981.
- [693] J. Smith and D. Smith. Database abstractions: Aggregation and generalization. *ACM Transactions on Database Systems*, 1(1):105–133, 1977.
- [694] K. Smith and M. Winslett. Entity modeling in the MLS relational model. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [695] P. Smith and M. Barnes. *Files and Databases: An Introduction*. Addison-Wesley, 1987.
- [696] N. Soparkar, H. Korth, and A. Silberschatz. Databases with deadline and contingency constraints. *IEEE Transactions on Knowledge and Data Engineering*, 7(4):552–565, 1995.
- [697] S. Spaccapietra, C. Parent, and Y. Dupont. Model independent assertions for integration of heterogeneous schemas. In *Proc. Intl. Conf. on Very Large Databases*, 1992.
- [698] S. Spaccapietra (ed.). *Entity-Relationship Approach: Ten Years of Experience in Information Modeling*, *Proc. Entity-Relationship Conf.* North-Holland, 1987.
- [699] E. Spertus. ParaSite: mining structural information on the web. In *Intl. World Wide Web Conference*, 1997.
- [700] R. Srikant and R. Agrawal. Mining generalized association rules. In *Proc. Intl. Conf. on Very Large Databases*, 1995.
- [701] R. Srikant and R. Agrawal. Mining Quantitative Association Rules in Large Relational Tables. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1996.
- [702] R. Srikant and R. Agrawal. Mining Sequential Patterns: Generalizations and Performance Improvements. In *Proc. Intl. Conf. on Extending Database Technology*, 1996.

- [703] R. Srikant, Q. Vu, and R. Agrawal. Mining Association Rules with Item Constraints. In *Proc. Intl. Conf. on Knowledge Discovery in Databases and Data Mining*, 1997.
- [704] V. Srinivasan and M. Carey. Performance of B-Tree concurrency control algorithms. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1991.
- [705] D. Srivastava, S. Dar, H. Jagadish, and A. Levy. Answering queries with aggregation using views. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [706] D. Srivastava, R. Ramakrishnan, P. Seshadri, and S. Sudarshan. Coral++: Adding object-orientation to a logic database language. In *Proc. Intl. Conf. on Very Large Databases*, 1993.
- [707] J. Srivastava and D. Rotem. Analytical modeling of materialized view maintenance. In *ACM Symp. on Principles of Database Systems*, 1988.
- [708] J. Srivastava, J. Tan, and V. Lum. Tbsam: An access method for efficient processing of statistical queries. *IEEE Transactions on Knowledge and Data Engineering*, 1(4):414–423, 1989.
- [709] D. Stacey. Replication: DB2 , Oracle or Sybase? *Database Programming and Design*, pages 42–50, December 1994.
- [710] P. Stachour and B. Thuraisingham. Design of LDV: A multilevel secure relational database management system. *IEEE Transactions on Knowledge and Data Engineering*, 2(2), 1990.
- [711] J. Stankovic and W. Zhao. On real-time transactions. In *Proc. ACM SIGMOD Conf. on the Management of Data Record*, 1988.
- [712] T. Steel. Interim report of the ANSI-SPARC study group. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1975.
- [713] M. Stonebraker. Implementation of integrity constraints and views by query modification. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1975.
- [714] M. Stonebraker. Concurrency control and consistency of multiple copies of data in Distributed Ingres. *IEEE Transactions on Software Engineering*, 5(3), 1979.
- [715] M. Stonebraker. Operating system support for database management. *Communications of the ACM*, 14(7):412–418, 1981.
- [716] M. Stonebraker. Inclusion of new types in relational database systems. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1986.

- [717] M. Stonebraker. *The INGRES Papers: Anatomy of a Relational Database System*. Addison-Wesley, 1986.
- [718] M. Stonebraker. The design of the Postgres storage system. In *Proc. Intl. Conf. on Very Large Databases*, 1987.
- [719] M. Stonebraker. *Object-relational DBMSs—The Next Great Wave*. Morgan Kaufmann, 1996.
- [720] M. Stonebraker, J. Frew, K. Gardels, and J. Meredith. The Sequoia 2000 storage benchmark. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.
- [721] M. Stonebraker and J. Hellerstein (eds). *Readings in Database Systems*. Morgan Kaufmann, 2 edition, 1994.
- [722] M. Stonebraker, A. Jhingran, J. Goh, and S. Potamianos. On rules, procedures, caching and views in data base systems. In *UCBERL M9036*, 1990.
- [723] M. Stonebraker and G. Kemnitz. The Postgres next-generation database management system. *Communications of the ACM*, 34(10):78–92, 1991.
- [724] B. Subramanian, T. Leung, S. Vandenberg, and S. Zdonik. The AQUA approach to querying lists and trees in object-oriented databases. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1995.
- [725] W. Sun, Y. Ling, N. Rishe, and Y. Deng. An instant and accurate size estimation method for joins and selections in a retrieval-intensive environment. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.
- [726] A. Swami and A. Gupta. Optimization of large join queries: Combining heuristics and combinatorial techniques. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1989.
- [727] T. Swift and D. Warren. An abstract machine for SLG resolution: Definite programs. In *Intl. Logic Programming Symposium*, 1994.
- [728] A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, and R. Snodgrass. *Temporal Databases: Theory, Design and Implementation*. Benjamin-Cummings, 1993.
- [729] Y. Tay, N. Goodman, and R. Suri. Locking performance in centralized databases. *ACM Transactions on Database Systems*, 10(4):415–462, 1985.
- [730] T. Teorey. *Database Modeling and Design: The E-R Approach*. Morgan Kaufmann, 1990.

- [731] T. Teorey, D.-Q. Yang, and J. Fry. A logical database design methodology for relational databases using the extended entity-relationship model. *ACM Computing Surveys*, 18(2):197–222, 1986.
- [732] R. Thomas. A majority consensus approach to concurrency control for multiple copy databases. *ACM Transactions on Database Systems*, 4(2):180–209, 1979.
- [733] S. A. Thomas. *SSL & TLS Essentials: Securing the Web*. John Wiley & Sons, 2000.
- [734] A. Thomasian. Concurrency control: Methods, performance, and analysis. *ACM Computing Surveys*, 30(1):70–119, 1998.
- [735] A. Thomasian. Two-phase locking performance and its thrashing behavior. *ACM Computing Surveys*, 30(1):70–119, 1998.
- [736] S. Thomas, S. Bodagala, K. Alsabti, and S. Ranka. An efficient algorithm for the incremental updation of association rules in large databases. In *Proc. Intl. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1997.
- [737] S. Todd. The Peterlee relational test vehicle. *IBM Systems Journal*, 15(4):285–307, 1976.
- [738] H. Toivonen. Sampling large databases for association rules. In *Proc. Intl. Conf. on Very Large Databases*, 1996.
- [739] TP Performance Council. TPC Benchmark D: Standard specification, rev. 1.2. Technical report, <http://www.tpc.org/dspec.html>, 1996.
- [740] I. Traiger, J. Gray, C. Galtieri, and B. Lindsay. Transactions and consistency in distributed database systems. *ACM Transactions on Database Systems*, 25(9), 1982.
- [741] M. Tsangaris and J. Naughton. On the performance of object clustering techniques. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1992.
- [742] D.-M. Tsou and P. Fischer. Decomposition of a relation scheme into Boyce-C odd normal form. *SIGACT News*, 14(3):23–29, 1982.
- [743] D. Tsur, J. D. Ullman, S. Abiteboul, C. Clifton, R. Motwani, S. Nestorov, and A. Rosenthal. Query flocks: A generalization of association-rule mining. In *Proc. ACM SIGMOD Conf. on Management of Data*, pages 1–12, 1998.
- [744] A. Tucker (ed.). *Computer Science and Engineering Handbook*. CRC Press, 1996.
- [745] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

- [746] J. Ullman. The U.R. strikes back. In *ACM Symp. on Principles of Database Systems*, 1982.
- [747] J. Ullman. *Principles of Database and Knowledgebase Systems, Vols. 1 and 2*. Computer Science Press, 1989.
- [748] J. Ullman. Information integration using logical views. In *Intl. Conf. on Database Theory*, 1997.
- [749] S. Urban and L. Delcambre. An analysis of the structural, dynamic, and temporal aspects of semantic data models. In *Proc. IEEE Intl. Conf. on Data Engineering*, 1986.
- [750] G. Valentin, M. Zuliani, D. C. Zilio, G. M. Lohman, and A. Skelley. Db2 advisor: An optimizer smart enough to recommend its own indexes. In *Proc. Intl. Conf. on Data Engineering (ICDE)*, pages 101–110. IEEE Computer Society, 2000.
- [751] M. Van Emden and R. Kowalski. The semantics of predicate logic as a programming language. *Journal of the ACM*, 23(4):733–742, 1976.
- [752] A. Van Gelder. Negation as failure using tight derivations for general logic programs. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufmann, 1988.
- [753] C. J. van Rijsbergen. *Information Retrieval*. Butterworths, London, United Kingdom, 1990.
- [754] M. Vardi. Incomplete information and default reasoning. In *ACM Symp. on Principles of Database Systems*, 1986.
- [755] M. Vardi. Fundamentals of dependency theory. In *Trends in Theoretical Computer Science*. ed. E. Borger, Computer Science Press, 1987.
- [756] L. Vieille. Recursive axioms in deductive databases: The query-subquery approach. In *Intl. Conf. on Expert Database Systems*, 1986.
- [757] L. Vieille. From QSQ towards QoSAQ: global optimization of recursive queries. In *Intl. Conf. on Expert Database Systems*, 1988.
- [758] L. Vieille, P. Bayer, V. Kuchenhoff, and A. Lefebvre. EKS-V1 , a short overview. In *AAAI-90 Workshop on Knowledge Base Management Systems*, 1990.
- [759] J. S. Vitter and M. Wang. Approximate computation of multidimensional aggregates of sparse data using wavelets. In *Proc. ACM SIGMOD Conf. on the Management of Data*, pages 193–204. ACM Press, 1999.
- [760] G. von Bultzingsloewen. Translating and optimizing SQL queries having aggregates. In *Proc. Intl. Conf. on Very Large Databases*, 1987.

- [761] G. von Bultzingsloewen, K. Dittrich, C. Iochpe, R.-P. Liedtke, P. Lockemann, and M. Schryro. Kardamom—A dataflow database machine for real-time applications. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1988.
- [762] G. Vossen. *Data Models, Database Languages and Database Management Systems*. Addison-Wesley, 1991.
- [763] N. Wade. Citation analysis: A new tool for science administrators. *Science*, 188(4183):429–432, 1975.
- [764] R. Wagner. Indexing design considerations. *IBM Systems Journal*, 12(4):351–367, 1973.
- [765] X. Wang, S. Jajodia, and V. Subrahmanian. Temporal modules: An approach toward federated temporal databases. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1993.
- [766] K. Wang and H. Liu. Schema discovery for semistructured data. In *Third International Conference on Knowledge Discovery and Data Mining (KDD -97)*, pages 271–274, 1997.
- [767] R. Weber, H. Schek, and S. Blott. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. In *Proc. Intl. Conf. on Very Large Data Bases*, 1998.
- [768] G. Weddell. Reasoning about functional dependencies generalized for semantic data models. *ACM Transactions on Database Systems*, 17(1), 1992.
- [769] W. Weihl. The impact of recovery on concurrency control. In *ACM Symp. on Principles of Database Systems*, 1989.
- [770] G. Weikum and G. Vossen. *Transactional Information Systems*. Morgan Kaufmann, 2001.
- [771] R. Weiss, B. V. Iez, M. A. Sheldon, C. Manprempre, P. Szilagyi, A. Duda, and D. K. Gifford. HyPursuit: A hierarchical network search engine that exploits content-link hypertext clustering. In *Proc. ACM Conf. on Hypertext*, 1996.
- [772] C. White. Let the replication battle begin. In *Database Programming and Design*, pages 21–24, May 1994.
- [773] S. White, M. Fisher, R. Cattell, G. Hamilton, and M. Hapner. *JDBC API Tutorial and Reference: Universal Data Access for the Java 2 Platform*. Addison-Wesley, 2 edition, 1999.
- [774] J. Widom and S. Ceri. *Active Database Systems*. Morgan Kaufmann, 1996.

- [775] G. Wiederhold. *Database Design (2nd ed.)*. McGraw-Hill, 1983.
- [776] G. Wiederhold, S. Kaplan, and D. Sagalowicz. Physical database design research at Stanford. *IEEE Database Engineering*, 1:117–119, 1983.
- [777] R. Williams, D. Daniels, L. Haas, G. Lapis, B. Lindsay, P. Ng, R. Obermarck, P. Selinger, A. Walker, P. Wilms, and R. Yost. R\*: An overview of the architecture. Technical report, IBM RJ3325, San Jose, CA, 1981.
- [778] M. S. Winslett. A model-based approach to updating databases with incomplete information. *ACM Transactions on Database Systems*, 13(2):167–196, 1988.
- [779] G. Wiorowski and D. Kull. *DB2: Design and Development Guide (3rd ed.)*. Addison-Wesley, 1992.
- [780] I. H. Witten, A. Moffat, and T. C. Bell. *Managing Gigabytes: Compressing and Indexing Documents and Images*. Van Nostrand Reinhold, 1994.
- [781] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann Publishers, 1999.
- [782] O. Wolfson, A. Sistla, , B. Xu, J. Zhou, and S. Chamberlain. Domino: Databases for moving objects tracking. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, 1999.
- [783] Y. Yang and R. Miller. Association rules over interval data. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1997.
- [784] K. Youssefi and E. Wong. Query processing in a relational database management system. In *Proc. Intl. Conf. on Very Large Databases*, 1979.
- [785] C. Yu and C. Chang. Distributed query processing. *ACM Computing Surveys*, 16(4):399–433, 1984.
- [786] O. R. Zaiane, M. El-Hajj, and P. Lu. Fast Parallel Association Rule Mining Without Candidacy Generation. In *Proc. IEEE Intl. Conf. on Data Mining (ICDM)*, 2001.
- [787] M. J. Zaki. Scalable algorithms for association mining. In *IEEE Transactions on Knowledge and Data Engineering*, volume 12, pages 372–390, May/June 2000.
- [788] M. J. Zaki and C.-T. Ho, editors. *Large-Scale Parallel Data Mining*. Springer Verlag, 2000.
- [789] C. Zaniolo. Analysis and design of relational schemata. Technical report, Ph.D. Thesis, UCLA, TR UCLA-ENG-7669, 1976.

- [790] C. Zaniolo. Database relations with null values. *Journal of Computer and System Sciences*, 28(1):142–166, 1984.
- [791] C. Zaniolo. The database language GEM. In *Readings in Object-Oriented Databases*. eds. S.B. Zdonik and D. Maier, Morgan Kaufmann, 1990.
- [792] C. Zaniolo. Active database rules with transaction-conscious stable-model semantics. In *Intl. Conf. on Deductive and Object-Oriented Databases*, 1996.
- [793] C. Zaniolo, N. Arni, and K. Ong. Negation and aggregates in recursive rules: the LDL++ approach. In *Intl. Conf. on Deductive and Object-Oriented Databases*, 1993.
- [794] C. Zaniolo, S. Ceri, C. Faloutsos, R. Snodgrass, V. Subrahmanian, and R. Zicari. *Advanced Database Systems*. Morgan Kaufmann, 1997.
- [795] S. Zdonik, U. Cetintemel, M. Cherniack, C. Convey, S. Lee, G. Seidman, M. Stonebraker, N. Tatbul, and D. Carney. Monitoring streams—A new class of data management applications. In *Proc. Intl. Conf. on Very Large Data Bases*, 2002.
- [796] S. Zdonik and D. Maier (eds.). *Readings in Object-Oriented Databases*. Morgan Kaufmann, 1990.
- [797] A. Zhang, M. Nodine, B. Bhargava, and O. Bukhres. Ensuring relaxed atomicity for flexible transactions in multidatabase systems. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1994.
- [798] T. Zhang, R. Ramakrishnan, and M. Livny. BIRCH: an efficient data clustering method for very large databases. In *Proc. ACM SIGMOD Conf. on Management of Data*, 1996.
- [799] Y. Zhao, P. Deshpande, J. F. Naughton, and A. Shukla. Simultaneous optimization and evaluation of multiple dimensional queries. In *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 1998.
- [800] Y. Zhuge, H. Garcia-Molina, J. Hammer, and J. Widom. View maintenance in a warehousing environment. In *Proc. ACM SIGMOD Conf. on the Management of Data*, 1995.
- [801] M. M. Zloof. Query-by-example: a database language. *IBM Systems Journal*, 16(4):324–343, 1977.
- [802] J. Zobel, A. Moffat, and K. Ramamohanarao. Inverted files versus signature files for text indexing. *ACM Transactions on Database Systems*, 23, 1998.
- [803] J. Zobel, A. Moffat, and R. Sacks-Davis. An efficient indexing technique for full text databases. In *Proc. Intl. Conf. on Very Large Databases, Morgan Kaufman pubs. (San Francisco, CA) 18, Vancouver*, 1992.

- [804] U. Zukowski and B. Freitag. The deductive database system LOLA. In *Proc. Intl. Conf. on Logic Programming and Non-Monotonic Reasoning*, 1997.