

Bibliography

- [Anahory and Murray, 1997] Anahory, S. and Murray, D. (1997). *Data Warehousing in the Real World*. Addison-Wesley, Essex.
- [Arge, 1995] Arge, L. (1995). The buffer tree: A new technique for optimal I/O algorithms. In *International Workshop on Algorithms and Data Structures*, LNCS 1460, pages 334–345, Berlin/New York. Springer.
- [Bayer, 1996] Bayer, R. (1996). The universal B-tree for multidimensional indexing. Technical Report TUM-I9637, Technische Universität München.
- [Bayer and Markl, 1998] Bayer, R. and Markl, V. (1998). The UB-tree: Performance of multidimensional range queries. Technical Report TUM-I9814, Technische Universität München.
- [Bayer and McCreight, 1972] Bayer, R. and McCreight, E. (1972). Organization and maintenance of large ordered indexes. *Acta Informatica*, 1(3):173–189.
- [Beckmann et al., 1990] Beckmann, N., Kriegel, H.-P., Schneider, R., and Seeger, B. (1990). The R*-tree: An efficient and robust access method for points and rectangles. In *Proceedings of the 1985 ACM SIGMOD International Conference on Management of Data*, pages 322–331, New York.
- [Bentley, 1975] Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517.
- [Berchtold et al., 1996] Berchtold, S., Keim, D. A., and Kriegel, H.-P. (1996). The X-tree: An index structure for high-dimensional data. In *Proceedings of the 22nd International Conference on Very Large Databases (VLDB)*, pages 28–39.
- [Bitton and Gray, 1998] Bitton, D. and Gray, J. (1998). The rebirth of database machines. Invited talk at the 24th International Conference on Very Large Data Bases (VLDB).
- [Box and Muller, 1958] Box, G. E. P. and Muller, M. E. (1958). A note on the generation of random normal deviates. *Ann. Math. Stat.*, 29:610–611.
- [Breimann et al., 1984] Breimann, L., Friedmann, J. H., Olshen, R. A., and Stone, C. J. (1984). *Classification and Regression Trees*. Wadsworth and Brooks/Cole, Monterey.

- [Chan and Ioannidis, 1998] Chan, C.-Y. and Ioannidis, Y. E. (1998). Bitmap index design and evaluation. In *Proceedings of the International Conference on Management of Data*, pages 355 – 366.
- [Chan and Ioannidis, 1999] Chan, C.-Y. and Ioannidis, Y. E. (1999). An efficient bitmap encoding scheme for selection queries. In *SIGMOD Conference 1999*, pages 215–226.
- [Chaudhuri and Dayal, 1997] Chaudhuri, S. and Dayal, U. (1997). An overview of data warehousing and OLAP technology. *SIGMOD Record*, 26(1):65–74.
- [Chen et al., 1998] Chen, L., Choubey, R., and Rundensteiner, E. A. (1998). Bulk-insertions into R-trees using the small-tree-large-tree approach. In *Proceedings of the 6th international symposium on Advances in geographic information systems*, pages 161–162.
- [Christiansen et al., 1998] Christiansen, A., Höding, M., Rautenstrauch, C., and Saake, G. (1998). *Oracle8 effizient einsetzen - Aufbau, Entwicklung, Verteilung und Betrieb leistungsfähiger Oracle8-Anwendungen*. Addison Wesley.
- [Codd, 1994] Codd, E. F. (1994). Adding value to relational and legacy DBMS: The OLAP mandate. *Business Intelligence*.
- [Faloutsos and Kamel, 1994] Faloutsos, C. and Kamel, I. (1994). Beyond uniformity and independence: Analysis of R- trees using the concept of fractal dimension. In *PODS '94. Proceedings of the Thirteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, volume 13, pages 4–13, New York. ACM Press.
- [Finkel and Bentley, 1974] Finkel, R. A. and Bentley, J. L. (1974). Quad trees: a data structure for retrieval on composite keys. *Acta Informatica*, 4(1):1–9.
- [Gaede and Günther, 1998] Gaede, V. and Günther, O. (1998). Multidimensional access methods. *ACM Computing Surveys*, 30(2):170–231.
- [Garcia-Molina et al., 1999] Garcia-Molina, H., Ullman, J. D., and Widom, J. (1999). *Database Implementation*. Prentice Hall, Upper Saddle River, New Jersey.
- [Graefe, 1993] Graefe, G. (1993). Query evaluation techniques for large databases. *ACM Computing Surveys*, 25(2):73–170.
- [Gray et al., 1997] Gray, J., Chaudhuri, S., Bosworth, A., Layman, A., Reichart, D., and Venkatrao, M. (1997). Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals. *Data Mining and Knowledge Discovery*, 1(1):29–53.
- [Gray and Reuter, 1993] Gray, J. and Reuter, A. (1993). *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann, San Mateo.

- [Großer, 1997] Großer, P. (1997). Verwaltung relationaler Daten mit Hilfe einer mehrdimensionalen Indexstruktur. Master's thesis, Humboldt-Universität zu Berlin, Institut für Informatik.
- [Günther, 1989] Günther, O. (1989). The design of the cell tree: An object-oriented index structure for geometric databases. In *Proceedings of the Fifth International Conference on Data Engineering (ICDE)*, pages 598–605. IEEE Computer Society.
- [Günther et al., 1998] Günther, O., Oria, V., Picouet, P., Saglio, J.-M., and Scholl, M. (1998). Benchmarking spatial joins *à la carte*. In *Proceedings of the 10th International Conference on Scientific and Statistical Database Management, Proceedings*, pages 32–41, New York. IEEE Computer Society.
- [Gupta et al., 1997] Gupta, H., Harinaryan, V., Rajaraman, A., and Ullman, J. D. (1997). Index selection for OLAP. In *Proceedings of the International Conference on Data Engineering (ICDE)*, pages 208–219.
- [Gurret and Rigaux, 1998] Gurret, C. and Rigaux, P. (1998). An integrated platform for the evaluation of spatial query processing strategies. In *Proceedings of the 9th International Conference on Database and Expert Systems Applications (DEXA)*, LNCS 1460, pages 757–766. Springer.
- [Guttman, 1984] Guttman, A. (1984). R-trees: A dynamic index structure for spatial searching. In *SIGMOD'84, Proceedings of Annual Meeting, Boston, Massachusetts*, pages 47–57. ACM Press, New York.
- [Haas, 1999] Haas, P. J. (1999). Techniques for online exploration of large object-relational datasets. In *International Conference on Scientific and Statistical Database Management (SSDBM)*, pages 4–12. IEEE Computer Society.
- [Hannig, 1996] Hannig, U. (1996). *Data Warehouse und Management Informations Systeme*. Schäffer-Poeschel, Stuttgart.
- [Härder and Rahm, 1999] Härder, T. and Rahm, E. (1999). *Datenbanksysteme - Konzepte und Techniken der Implementierung*. Springer, Berlin, Heidelberg.
- [Härder and Reuter, 1983] Härder, T. and Reuter, A. (1983). Principles of transaction-oriented database recovery. *Computing Surveys*, 15(4):287–317.
- [Hellerstein et al., 1997a] Hellerstein, J. M., Haas, P. J., and Wang, H. J. (1997a). Online aggregation. In *SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data*, pages 171–182. ACM Press, New York.
- [Hellerstein et al., 1997b] Hellerstein, J. M., Koutsoupias, E., and Papadimitriou, C. H. (1997b). On the analysis of indexing schemes. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 249–256. ACM Press, New York.

- [Hellerstein et al., 1995] Hellerstein, J. M., Naughton, J. F., and Pfeffer, A. (1995). Generalized search trees for database systems. In *Proceedings of 21th International Conference on Very Large Data Bases (VLDB)*, pages 562–573. Morgan Kaufmann.
- [Ho et al., 1997] Ho, C.-T., Agrawal, R., Megiddo, N., and Srikant, R. (1997). Range queries in OLAP data cubes. In Peckham, J., editor, *SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data*, pages 73–88. ACM Press, New York.
- [Huber, 1981] Huber, P. J. (1981). *Robust Statistics*. Wiley, New York.
- [Huyn, 1997] Huyn, N. (1997). Multiple-view self-maintenance in data warehousing environments. In *Proceedings of the 23rd International Conference on Very Large Databases (VLDB)*, pages 26–35. Morgan Kaufmann, San Francisco, CA.
- [Ibbetson, 1963] Ibbetson, D. (1963). Collected algorithms. *Commun. Ass. Computing Mach.*, 6:616. Algorithm 209 Gauss.
- [Informix, 1997] Informix (1997). Indexing for the enterprise data warehouse. White paper. Available at <http://www.informix.com>.
- [Inmon, 1996] Inmon, W. H. (1996). *Building the Data Warehouse*. Wiley & Sons, New York.
- [Inmon et al., 1997] Inmon, W. H., Welch, J. D., and Katherine, G. L. (1997). *Managing the Data Warehouse*. Wiley & Sons, New York.
- [Jürgens and Lenz, 1998] Jürgens, M. and Lenz, H.-J. (1998). The R_a^* -tree: An improved R^* -tree with materialized data for supporting range queries on OLAP-data. In *Proceedings of the International Workshop on Data Warehouse Design and OLAP Technology (DWDOT98)*, pages 186–191. IEEE Computer Society Press.
- [Jürgens and Lenz, 1999a] Jürgens, M. and Lenz, H.-J. (1999a). PISA: Performance models of index structures with and without aggregated data. In *Proceedings of the 8th International Conference on Statistical and Scientific Database Management (SS-DBM)*, pages 78–87. IEEE Computer Society.
- [Jürgens and Lenz, 1999b] Jürgens, M. and Lenz, H.-J. (1999b). Tree based indexes vs. bitmap indexes: A performance study. In *International Workshop on Design and Management of Data Warehouses, Heidelberg*, pages 78–87.
- [Kamel and Faloutsos, 1993] Kamel, I. and Faloutsos, C. (1993). On packing R-trees. In *2nd International Conference on Information and Knowledge Management (CIKM)*, pages 490–499. ACM Press, New York.
- [Kamel and Faloutsos, 1994] Kamel, I. and Faloutsos, C. (1994). Hilbert R-tree: An improved R-tree using fractals. In *Proceedings of the 20th International Conference on Very Large Data Bases (VLDB)*, pages 500–509.

- [Karayama and Satoh, 1997] Karayama, N. and Satoh, S. (1997). The SR-tree: An index structure for high-dimensional nearest neighbour queries. In *Proceedings ACM SIGMOD International Conference on Management of Data*, pages 369–380.
- [Kornacker, 1999] Kornacker, M. (1999). High-performance extensible indexing. In *Proceedings of the 25rd International Conference on Very Large Databaseses (VLDB)*, pages 699–708. Morgan Kaufmann, San Farnsisco, CA.
- [Kornacker et al., 1997] Kornacker, M., Mohan, C., and Hellerstein, J. M. (1997). Concurrency and recovery in generalized search trees. In Peckham, J., editor, *SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data*, pages 62–72. ACM Press, New York.
- [Kornacker et al., 1998] Kornacker, M., Shah, M., and Hellerstein, J. M. (1998). amdb: an access method debugging tool. In *SIGMOD 1998, Proceedings ACM SIGMOD International Conference on Management of Data, June 2-4, 1998*, pages 570–571. ACM Press, New York.
- [Kuan and Lewis, 1999] Kuan, J. and Lewis, P. (1999). A study on data point search for HG-trees. *SIGMOD Record*, 28(1):90–96.
- [Labio et al., 1997] Labio, W. J., Quass, D., and Adelberg, B. (1997). Physical database design for data warehouses. In *Proceedings of the ICDE*, pages 277–288.
- [Lamersdorf et al., 1996] Lamersdorf, W., Lenz, H.-J., and Rieger, B. (1996). *Data Warehousing, OLAP, Führungsinformationssysteme... Neue Entwicklungen des Informationsmanagements*. Congress VIII, ONLINE GmbH Kongresse und Messen für Technische Kommunikation, Velbert.
- [Lee and Hammer, 1999] Lee, M. and Hammer, J. (1999). Speeding up warehouse physical desing using a randomized algorithm. In *Proceedings of the International Workshop on Design and Management of Data Warehouses (DMDW'99)*.
- [Lehner et al., 1998] Lehner, W., Albrecht, J., and Wedekind, H. (1998). Normal forms for multidimensional databases. In *10th International Conference on Scientific and Statistical Database Management*, pages 63–72. IEEE Computer Society.
- [Lenz, 1993] Lenz, H.-J. (1993). On the design of statistical databases, micro-, marco- und meta-database modelling. In Faulbaum, F., editor, *Advances in Statistical Software 5*. Gustav Fischer, Stuttgart.
- [Lenz and Jürgens, 1998] Lenz, H.-J. and Jürgens, M. (1998). Modeling and improving the performance of multidimensional indexstructures for range queries on OLAP data. Technical Report 1998/29, Fachbereich Wirtschaftswissenschaft der Freien Universität Berlin.

- [Lenz and Shoshani, 1997] Lenz, H. J. and Shoshani, A. (1997). Summarizability in OLAP and statistical data bases. In *Proceedings of 9th International Conference on Statistical and Scientific Database Management*, pages 132–143. IEEE Computer Society Press.
- [Leutenegger and Lopez, 1998] Leutenegger, S. T. and Lopez, M. A. (1998). The effect of buffering on the performance of R-trees. In *Proceedings of the 14th International Conference on Data Engineering (ICDE)*, pages 164–171. IEEE Computer Society Press.
- [Leutenegger et al., 1997] Leutenegger, S. T., López, M. A., and Edgington, J. M. (1997). STR: A simple and efficient algorithm for R-tree packing. In *Proceedings of the 13th International Conference on Data Engineering (ICDE)*, pages 497–506, Los Alamitos, California. IEEE Computer Society Press.
- [Markl, 1999] Markl, V. (1999). *MISTRAL: Processing Relational Queries using a Multi-dimensional Access Technique*. PhD thesis, Technische Universität München.
- [Marques et al., 1998] Marques, P., Furtado, P., and Baumann, P. (1998). An efficient strategy for tiling multidimensional OLAP data cubes. In *GI Workshop on Data Mining and Data Warehousing, Magdeburg*, pages 13–24.
- [O’Neil and Quass, 1997] O’Neil, P. and Quass, D. (1997). Improved query performance with variant indexes. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 26(2):38–49.
- [Pagel, 1995] Pagel, B.-U. (1995). *Analyse und Optimierung von Indexstrukturen in Geodatenbanksysteme*. PhD thesis, Fernuniversität Hagen.
- [Pagel et al., 1993] Pagel, B.-U., Six, H.-W., Toben, H., and Widmayer, P. (1993). Towards an analysis of range query performance in spatial data structures. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 214–221. ACM Press, New York.
- [Patterson and Keeton, 1998] Patterson, D. A. and Keeton, K. K. (1998). Hardware technology trends and database opportunities. invited talk at SIGMOD 1998 International Conference on Management of Data.
- [Robinson, 1981] Robinson, J. T. (1981). The K-D-B-tree: A search structure for large multidimensional dynamic indexes. In *Proceedings of the ACM SIGMOD*, pages 10–18.
- [Rotem and Zhao, 1996] Rotem, D. and Zhao, J. L. (1996). Extendible arrays for statistical databases and OLAP applications. In *Proceedings of the 8th International Conference on Statistical and Scientific Database Management (SSDBM)*, pages 108–117. IEEE Computer Society.

- [Roussopoulos and Leifker, 1985] Roussopoulos, N. and Leifker, D. (1985). Direct spatial search on pictorial databases using packed R-trees. In *ACM SIGMOD (International Conference on Management of Data)*, pages 17–31, Austin, Texas.
- [Saake and Heuer, 1999] Saake, G. and Heuer, A. (1999). *Datenbanken: Implementierungstechniken*. MITP-Verlag, Bonn.
- [Samet, 1989] Samet, H. (1989). *Applications of Spatial Data Structures*. Addison-Wesley, Reading, Massachusetts.
- [Samet, 1990] Samet, H. (1990). *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading, Massachusetts.
- [Sapia, 1999] Sapia, C. (1999). On modeling and predicting user behaviour in OLAP systems. In *Proceedings of the International Workshop on Design and Management of Data Warehouses (DMDW'99)*.
- [Sarawagi, 1997] Sarawagi, S. (1997). Indexing OLAP data. *Data Engineering Bulletin*, 20(1):36–43.
- [Schnitzer and Leutenegger, 1999] Schnitzer, B. and Leutenegger, S. T. (1999). Master-client R-trees: A new parallel R-tree architecture. In *International Conference on Scientific and Statistical Database Management (SSDBM)*, pages 68–77.
- [Sellis et al., 1985] Sellis, T., Roussopoulos, N., and Faloutsos, C. (1985). The R^+ -tree: A dynamic index for multi-dimensional objects. In *Proceedings of the 13th international Conference on Very Large Databases (VLDB)*, pages 507–518.
- [Shoshani, 1997] Shoshani, A. (1997). OLAP and statistical databases: Similarities and differences. In *Proceedings of the 16th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 185–196.
- [Shukla et al., 1996] Shukla, A., Deshpande, P. M., Naughton, J. F., and Ramasamy, K. (1996). Storage estimation for multidimensional aggregates in the presence of hierarcies. In *Proceedings of the 22nd International Conference on Very Large Databases (VLDB)*, pages 522–531.
- [Srivastava et al., 1989] Srivastava, J., Tan, J. S. E., and Lum, V. Y. (1989). TBSAM: An access method for efficient processing of statistical queries. *IEEE Transactions on Knowledge and Data Engineering*, 1(4):414–423.
- [Suhl, 1998] Suhl, U. H. (1998). Mathematical optimizing system. Further information is available at <http://mops.wiwiss.fu-berlin.de>.
- [Sybase, 1997] Sybase (1997). Adaptive server IQ. White paper. available at <http://www.sybase.com>.

- [Theodoridis and Sellis, 1996] Theodoridis, Y. and Sellis, T. K. (1996). A model for the prediction of R-tree performance. In *Proceedings of the Fifteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 161–171. ACM Press, New York.
- [van den Bercken et al., 1997] van den Bercken, J., Seeger, B., and Widmayer, P. (1997). A generic approach to bulk loading multidimensional index structures. In *Proceedings of the 23rd International Conference on Very Large Databases (VLDB)*, pages 406–415.
- [Venables and Ripley, 1994] Venables, W. N. and Ripley, B. D. (1994). *Modern Applied Statistics with S-Plus*. Springer Verlag, New York, Berlin.
- [White and Jain, 1996] White, D. A. and Jain, R. (1996). Similarity indexing with the SS-tree. In *Proceedings of the 12th International Conference on Data Engineering, New Orleans*, pages 516–523.
- [Wu and Buchmann, 1998] Wu, M.-C. and Buchmann, A. P. (1998). Encoded bitmap indexing for data warehouses. In *Proceedings of the 14th International Conference on Data Engineering (ICDE)*, pages 220–230.
- [Yang et al., 1997] Yang, J., Karlapalem, K., and Li, Q. (1997). Algorithms for materialized view design in data warehousing environment. In *Proceedings of the 23rd International Conference on very Large Databases (VLDB)*, pages 136–145.