Downloaded from UvA-DARE, the institutional repository of the University of Amsterdam (UvA)
http://dare.uva.nl/document/83751

File ID	83751
Filename	List of Publications

SOURCE (OR PART OF THE FOLLOWING SOURCE):

Type	Dissertation
Title	Principles of probabilistic query optimization
Author	F. Waas
Faculty	Faculty of Science
Year	2000

FULL BIBLIOGRAPHIC DETAILS:
http://dare.uva.nl/record/100085

Copyright

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other then for strictly personal, individual use.

UvA-DARE is a service provided by the library of the University of Amsterdam (http://dare.uva.nl)

List of Publications

[1] F. Waas and C. A. Galindo-Legaria. The Effect of Cost Distributions on Genetic Algorithms. Submitted for publication, October 2000.

[2] F. Waas and M. L. Kersten. Memory Aware Query Scheduling in a Database Cluster. Submitted for publication, June 2000.

[3] F. Waas and A. Pellenkoft. Join Order Selection — Good Enough is Easy. In *Proc. of the British National Conference on Databases*, Lecture Notes in Computer Science, pages 51-67, Exeter, United Kingdom, July 2000.

[4] A. R. Schmidt, M. L. Kersten, M. A. Windhouwer, and F. Waas. Efficient Relational Storage and Retrieval of XML Documents. In *International Workshop on the Web and Databases (In conjunction with ACM SIGMOD)*, pages 47-52, Dallas, TX, USA, May 2000.

[5] F. Waas and C. A. Galindo-Legaria. Counting, Enumerating and Sampling of Execution Plans in a Cost-Based Query Optimizer. In *Proc. of the ACM SIGMOD Int'l. Conf. on Management of Data*, pages 499-509, Dallas, TX, USA, May 2000.

[6] F. Waas. Extending Iterators for Advanced Query Execution. In *Australasian Database Conference*, pages 135-139, Canberra, Australia, January 2000. IEEE Computer Society Press.

[7] F. Waas and C. A. Galindo-Legaria. A Method for Counting, Enumerating and Sampling of Execution Plans in a Cost-Based Query Optimizer. MS-D 138320.1, SLWK-D 777.309US1, January 2000.

[8] F. Waas. Handling Non-deterministic Data Availability in Parallel Query Execution. In *Int'l. Workshop on Parallel and Distributed Databases*, pages 61-65, Florence, Italy, September 1999.

[9] S. Manegold and F. Waas. Integrating I/O processing and Transparent Parallelism—Toward Comprehensive Query Execution in Parallel Database Systems. In A. Dogac, M. T. Özsu, and O. Ulusoy, editors, *Current Trends in Database Systems*, pages 130-152. Idea Group Publisher, January 1999.

[10] S. Manegold, F. Waas, and M. L. Kersten. On Optimal Pipeline Processing in Parallel Query Optimization. In *Int'l. Conf. on Management of Data*, pages 217–235, Hyderabad, India, December 1998.

[11] S. Manegold and F. Waas. Thinking Big in a Small World—Efficient Parallel Query Execution on Small-scale SMPs. In *High Performance Computing Systems and Applications*, pages 133–146. Kluwer Academic Publishers, Dordrecht, The Netherlands, May 1998.

[12] S. Manegold, F. Waas, and D. Gudlat. *In Quest of the Bottleneck - Monitoring Parallel Database Systems*, volume 1332 of *Lecture Notes in Computer Science*, pages 277–284. Springer-Verlag, Berlin, New York, etc., November 1997.

[13] S. Manegold, J. K. Obermaier, and F. Waas. Load Balanced Query Evaluation in Shared-Everything Environments. In *Proc. of the European Conf. on Parallel Processing*, pages 1117–1124, Passau, Germany, August 1997.

[14] J. K. Obermaier and F. Waas. Dynamic Aspects of Query Processing in Parallel Database Systems. In *Proc. of the Workshop on Information Technologies and Systems*, pages 223–232, Cleveland, OH, USA, December 1996.