




SIGMOD 1997

Proceedings ACM SIGMOD
International Conference
on Management of Data

May 13-15, 1997
Tucson, Arizona, USA

Edited by Joan M. Peckman
SIGMOD Record, Volume 26, Issue 2, June 1997

Table of Contents

Concept and Experience Track

Session: Similarity Detection

Chair: Patrick Valduriez (INRIA)

Fast Parallel Similarity Search in Multimedia Databases

Stefan Berchtold, Christian Bohm, Bernhard Braunmuller
Daniel Keim, and Hans-Peter Kriegel 1

Similarity-Based Queries for Time Series Data

Davood Rafiei and Alberto Mendelzon 13

Meaningful Change Detection in Structured Data

Sudarshan S. Chawathe and Hector Garcia-Molina 26

Session: Index Management

Chair: Ron Morrison (University of St. Andrews)

Improved Query Performance with Variant Indexes

Patrick O'Neil and Dallan Quass 38

Highly Concurrent Cache Consistency for Indices in Client-Server Database Systems

Markos Zaharioudakis and Michael J. Carey 50

Concurrency and Recovery in Generalized Search Trees

Marcel Kornacker, C. Mohan, and Joseph M. Hellerstein 62

Session: Data Cubes and OLAP

Chair: Gail Mitchell (GTE)

Range Queries in OLAP Data Cubes

Ching-Tien Ho, Rakesh Agrawal, Nimrod Megiddo,
and Ramakrishnan Srikant 73

Cubetree: Organization of and Bulk Updates on the Data Cube

Nick Roussopoulos, Yannis Kotidis, and Mema Roussopoulos 89

Maintenance of Data Cubes and Summary Tables in a Warehouse

Inderpal Singh Mumick, Dallan Quass, and Barinderpal Singh Mumick 100

Session: Performance and Benchmarking

Chair: Louiga Raschid (University of Maryland)

Database Buffer Size Investigation for OLTP Workloads

Thin-Fong Tsuei, Allan N. Packer, and Keng-Tai Ko 112

<i>Database Performance in the Real World: TPC-D and SAP R/3</i> Joachen Doppelhammer, Thomas Hoppler, Alfons Kemper, and Donald Kossmann	123
<i>The BUCKY Object-Relational Benchmark</i> Michael J. Carey, David J. DeWitt, Jeffrey F. Naughton, Mohammad Asgarian, Paul Brown, Johannes E. Gehrke, and Dhaval N. Shah	135
Session: Aggregate and Derived Data	
Chair: H. V. Jagadish (AT&T Labs)	
<i>The STRIP Rule System For Efficiently Maintaining Derived Data</i> Brad Adelberg, Hector Garcia-Molina, and Jennifer Widom	147
<i>An Array-Based Algorithm for Simultaneous Multidimensional Aggregates</i> Yihong Zhao, Prasad M. Deshpande, and Jeffrey F. Naughton	159
<i>Online Aggregation</i> Joseph M. Hellerstein, Peter J. Haas, and Helen J. Wang	171
Session: Widely Distributed Data	
Chair: Waqar Hasan (Informix)	
<i>Balancing Push and Pull for Data Broadcast</i> Swarup Acharya, Michael Franklin, and Stanley Zdonik	183
<i>InfoSleuth: Agent-Based Semantic Integration of Information in Open and Dynamic Environments</i> R. J. Bayardo Jr., W. Bohrer, R. Brice, A. Cichocki, J. Fowler, A. Helal, V. Kashyap, T. Ksiezyk, G. Martin, M. Nodine, M. Rashid, M. Rusinkiewicz, R. Shea, C. Unnikrishnan, A. Unruh, and D. Woelk	195
<i>STARTS: Stanford Proposal for Internet Meta-Searching</i> Luis Gravano, Chen-Chuan K. Chang, Hector Garcia-Molina, and Andreas Paepcke	207
Session: Queries and Sorting	
Chair: Phil Bernstein (Microsoft)	
<i>On Saying "Enough Already!" in SQL</i> Michael J. Carey and Donald Kossmann	219
<i>A Framework for Implementing Hypothetical Queries</i> Timothy Griffin and Richard Hull	231

<i>High-Performance Sorting on Networks of Workstations</i> Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau David E. Culler, Joseph M. Hellerstein, and David A. Patterson	243
---	-----

Session: Data Mining

Chair: Peter Bunemann (University of Pennsylvania)

<i>Dynamic Itemset Counting and Implication Rules for Market Basket Data</i> Sergey Brin, Rajeev Motwani, Jeffrey D. Ullman, and Shalom Tsur	255
<i>Beyond Market Baskets: Generalizing Association Rules to Correlations</i> Sergey Brin, Rajeev Motwani, and Craig Silverstein	265
<i>Scalable Parallel Data Mining for Association Rules</i> Eui-Hong (Sam) Han, George Karypis, and Vipin Kumar	277

Session: Complex Data Access

Chair: Tiziana Catarci (Universita' degli Studi di Roma "La Sapienza")

<i>Efficiently Supporting Ad Hoc Queries in Large Datasets of Time Sequences</i> Flip Korn, H.V. Jagadish, and Christos Faloutsos	289
<i>DEVise: Integrated Querying and Visual Exploration of Large Datasets</i> M. Livny, R. Ramakrishnan, K. Beyer, G. Chen, D. Donjerkovic, S. Lawande, J. Myllymaki, and K. Wenger	301
<i>Partitioned Garbage Collection of a Large Object Store</i> Umesh Maheshwari and Barbara Liskov	313

Session: Spatial Data

Chair: Christos Faloutsos (University of Maryland)

<i>Size Separation Spatial Join</i> Nick Koudas and Kenneth C. Sevcik	324
<i>Building a Scaleable Geo-Spatial DBMS: Technology, Implementation, and Evaluation</i> Jignesh Patel, JieBing Yu, Navin Kabra, Kristin Tufte, Biswadeep Nag, Josef Burger, Nancy Hall, Karthikeyan Ramasamy, Roger Lueder, Curt Ellmann, Jim Kupsch, Shelly Guo, Johan Larson, David DeWitt, and Jeffrey Naughton	336
<i>A Toolkit for Negotiation Support Interfaces to Multi-Dimensional Data</i> Michael Gebhardt, Matthias Jarke, and Stephan Jacobs	348

Session: Novel Index Structures

Chair: Timos Sellis (National Technical University of Athens)

<i>Distance-Based Indexing for High-Dimensional Metric Spaces</i> Tolga Bozkaya and Meral Ozsoyoglu	357
--	-----

<i>The SR-tree: An Index Structure for High-Dimensional Nearest Neighbor Queries</i> Norio Katayama and Shin'ichi Satoh	369
--	-----

<i>Wave-Indices: Indexing Evolving Databases</i> Narayanan Shivakumar and Hector Garcia-Molina	381
---	-----

Session: View Maintenance

Chair: Narain Gehani (Bell Labs, Lucent Technologies)

<i>On-line Warehouse View Maintenance</i> Dallan Quass and Jennifer Widom	393
--	-----

<i>Supporting Multiple View Maintenance Policies</i> Latha S. Colby, Akira Kawaguchi, Daniel F. Lieuwen, Inderpal Singh Mumick, and Kenneth A. Ross	405
---	-----

<i>Efficient View Maintenance at Data Warehouses</i> D. Agrawal, A. El Abbadi, A. Singh, and T. Yurek	417
--	-----

Session: Rule Management

Chair: Isabel Cruz, Brown University

<i>Eliminating Costly Redundant Computations from SQL Trigger Executions</i> Francois Llirbat, Francoise Fabret, and Eric Simon	428
--	-----

<i>Temporal Aggregation in Active Database Rules</i> Iakovos Motakis and Carlo Zaniolo	440
---	-----

<i>Association Rules over Interval Data</i> R.J. Miller and Y. Yang	452
--	-----

Session: Security and Commit

Chair: Hank Korth (Bell Labs, Lucent Technologies)

<i>Secure Transaction Processing in Firm Real-Time Database Systems</i> Binto George and Jayant Haritsa	462
--	-----

<i>A Unified Framework for Enforcing Multiple Access Control Policies</i> Sushil Jajodia, Pierangela Samarati, V. S. Subrahmanian, and Eliza Bertino	474
--	-----

<i>Revisiting Commit Processing in Distributed Database Systems</i> Ramesh Gupta, Jayant Haritsa, and Krithi Ramamritham	486
---	-----

Tutorials

<i>Lessons from Wall Street</i> Dennis Shasha	498
<i>Object-Relational Database Systems: Principles, Products and Challenges</i> Michael J. Carey, Nelson M. Mattos, and Anil K. Nori	502
<i>Databases on the Web: Technologies for Federation Architectures and Case Studies</i> Ralf Kramer	503
<i>Data Warehousing and OLAP for Decision Support</i> Surajit Chaudhuri and Umeshwar Dayal	507

Panel

<i>Query Optimization at the Crossroads</i> Moderator: Surajit Chaudhuri	509
---	-----

Research Prototype Demonstrations

Visual Interfaces

<i>Delaunay: A Database Visualization System</i> Isabel F. Cruz, M. Averbuch, Wendy T. Lucas, and Melissa Radzysinski	510
<i>Picture Programming Project</i> Nita Goyal, Charles Hoch, Ravi Krishnamurthy, Brian Meckler, Michael Suchow, and Moshe Zloof	514
<i>DEVise: Integrated Querying and Visual Exploration of Large Datasets</i> M. Livny, R. Ramakrishnan, K. Beyer, G. Chen, D. Donjerkovic, S. Lawande, J. Myllymaki, and K. Wenger	517
<i>SEMCOG: An Object-Based Image Retrieval System and Its Visual Query Interface</i> Wen-Syan Li, K. Selcuk Candan, Kyoji Hirata, and Yoshinori Hara	521

Multi-Databases and Interoperability

<i>The COntext INterchange Mediator Prototype</i> S. Bressan, C.H. Goh, K. Fynn, M. Jakobisiak, K. Hussein, H. Kon, T. Lee, S. Madnick, T. Pena, J. Qu, A. Shum, and M. Siegel	525
---	-----

<i>MDM: A Multiple-Data Model Tool for the Management of Heterogeneous Database Schemes</i> Paolo Atzeni and Riccardo Torlone	528
<i>Template-Based Wrappers in the TSIMMIS System</i> Joachim Hammer, Hector Garcia-Molina, Svetlozar Nestorov, Ramana Yerneni, Marcus Breunig, and Vasilis Vassalos	532
<i>Languages for Multi-Database Interoperability</i> Frederic Gingras, Laks V.S. Lakshmanan, Iyer N. Subramanian, Despina Papoulis, and Nematollaah Shiri	536
 Accessing Data on the WWW	
<i>Infomaster: An Information Integration System</i> Michael R. Genesereth, Arthur M. Keller, and Oliver M. Duschka	539
<i>The InfoSleuth Project</i> R. J. Bayardo Jr., W. Bohrer, R. Brice, A. Cichocki, J. Fowler, A. Helal, V. Kashyap, T. Ksiezyk, G. Martin, M. Nodine, M. Rashid, M. Rusinkiewicz, R. Shea, C. Unnikrishnan, A. Unruh, and D. Woelk	543
<i>The Distributed Information Search Component (Disco) and the World Wide Web</i> Anthony Tomasic, Remy Amouroux, Philippe Bonnet, Olga Kapitskaia, Hubert Naacke, and Louiqa Raschid	546
<i>STRUDEL: A Web Site Management System</i> Mary Fernandez, Daniela Florescu, Jaewoo Kang, Alon Levy, and Dan Suciu	549
 Data Warehousing and Mining	
<i>GeoMiner: A System Prototype for Spatial Data Mining</i> Jaiwei Han, Krzysztof Koperski, and Nebojsa Stefanovic	553
<i>The WHIPS Prototype for Data Warehouse Creation and Maintenance</i> Wilburt J. Labio, Yue Zhuge, Janet L. Wiener, Himanshu Gupta, Hector Garcia-Molina, and Jennifer Widom	557
 Non-Standard DBMSs	
<i>Structural Matching and Discovery in Document Databases</i> Jason Tsong-Li Wang, Dennis Shasha, George J.S. Chang, Liam Relihan, Kaizhong Zhang, and Girish Patel	560
<i>S3: Similarity Search in CAD Database Systems</i> Stefan Berchtold and Hans-Peter Kriegel	564

<i>PREDATOR: An OR-DBMS with Enhanced Data Types</i> Praveen Seshadri and Mark Paskin	568
--	-----

<i>SENTINEL: An Object-Oriented DBMS With Event-Based Rules</i> S. Chakravarthy	572
--	-----

Environments

<i>The MENTOR Workbench for Enterprise-Wide Workflow Management</i> Dirk Wodtke, Jeanine Weissenfels, Gerhard Weikum, Angelika Kotz Dittrich, and Peter Muth	576
--	-----

<i>ZOO: A Desktop Experiment Management Environment</i> Yannis E. Ioannidis, Miron Livny, Anastassia Ailamaki, Anand Narayanan, and Andrew Therber	580
--	-----

Author Index	584
---------------------------	-----