

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

New York University, NY, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Yannis Ioannidis Marc H. Scholl
Joachim W. Schmidt Florian Matthes
Mike Hatzopoulos Klemens Boehm
Alfons Kemper Torsten Grust
Christian Boehm (Eds.)

Advances in Database Technology – EDBT 2006

10th International Conference on Extending Database Technology
Munich, Germany, March 26-31, 2006
Proceedings



Springer

Volume Editors

Yannis Ioannidis

E-mail: yannis@di.uoa.gr

Marc H. Scholl

E-mail: Marc.Scholl@uni-konstanz.de

Joachim W. Schmidt

E-mail: j.w.schmidt@tu-harburg.de

Florian Matthes

E-mail: matthes@in.tum.de

Mike Hatzopoulos

E-mail: mike@di.uoa.gr

Klemens Boehm

E-mail: boehm@ipd.uka.de

Alfons Kemper

E-mail: Alfons.Kemper@in.tum.de

Torsten Grust

E-mail: grust@in.tum.de

Christian Boehm

E-mail: boehm@dbs.ifi.lmu.de

Library of Congress Control Number: 2006922026

CR Subject Classification (1998): H.2, H.4, H.3, C.2.4, K.4.4

LNCS Sublibrary: SL 3 – Information Systems and Application, incl. Internet/Web and HCI

ISSN 0302-9743

ISBN-10 3-540-32960-9 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-32960-2 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006

Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11687238 06/3142 5 4 3 2 1 0

Preface

The series of International Conferences on Extending Database Technology (EDBT) is an established and prestigious forum for the exchange of the latest research results in data management. It provides unique opportunities for database researchers, practitioners, developers, and users to explore new ideas, techniques, and tools, and to exchange experiences. This volume contains the proceedings of the 10th EDBT Conference, held in Munich, Germany, March 27-29, 2006. The conference included 3 keynote talks, 56 full-size and 4 half-size research papers in 20 sessions, 8 industrial presentations in 3 sessions, 1 panel session, 5 tutorials in 7 sessions, and 20 demonstrations in 4 sessions. All of the research papers as well as papers and abstracts from most of the other sessions are included here.

Distinguished members of the database and information-retrieval communities delivered the three keynotes, which were all in the spirit of the banner theme chosen for EDBT 2006: “From Database Systems to Universal Data Management.” Martin Kersten, a pioneer in the area of database support for ambient application environments and the investigator of several kernel database architectures, discussed various hard issues that arise in organic database systems, i.e., systems that can be embedded in several hardware applications and have autonomic behavior. Alan Smeaton, a leader in content-based retrieval of information in a wide variety of media, introduced us to the world of digital video libraries and challenged us with several open problems associated with their effective management. Finally, David Maier, whose introduction of object-oriented concepts to the field has been pivotal for the establishment of object-relational databases as the current state of the art, used the Lewis and Clark expedition as an analogy to discuss the problems faced when trying to deploy dataspace systems, i.e., systems that manage enterprise data in the entire spectrum from fully structured to completely unstructured.

The Research Program Committee consisted of 76 members and was chaired by Yannis Ioannidis (University of Athens). It accepted 60 papers (56 regular-size and 4 half-size) out of 352 submissions, both of which are the largest numbers ever for EDBT. Papers were submitted from 37 countries. The reviewing process was managed by the Microsoft Conference Management Toolkit, developed and supported by Surajit Chaudhuri and the CMT team (from Microsoft Research).

The Industrial and Applications Program was assembled by a small committee under Alfons Kemper (Technical University of Munich). The 8 short papers that appear in the proceedings were selected from 28 submissions.

One panel session was solicited by Klemens Böhm (University of Karlsruhe), who also selected five tutorials out of ten submissions.

Christian Böhm (University of Munich) led a 14-member committee in putting together the Demonstrations Program, consisting of 20 demonstrations selected from 52 submitted proposals.

Mike Hatzopoulos (University of Athens) edited the proceedings. He also produced an electronic version for inclusion in the SIGMOD digital library and for posting on the Web prior to the conference.

The program and social activities of EDBT 2006 are the result of a huge effort by many hundreds of authors, reviewers, presenters, and organizers. We thank them all for helping to make the conference a success. In particular, we want to thank Gisela

Krügel and Josef Lankes (both from the Technical University of Munich) for the smooth local organization.

Finally, this 10th edition of the EDBT Conference series coincides with the year of retirement of Joachim Schmidt, who was one of the “fathers” of EDBT. The success of the EDBT Association, its conferences, and services to the community is largely due to the effort and expertise that its founders invested. As a tribute to Joachim’s commitment and dedication to EDBT, the organizers invited Joachim as an Honorary Chair of this year’s conference. Thank you Joachim for all these years!

Yannis Ioannidis	Program Chair
Marc H. Scholl	General Chair
Florian Matthes	Executive Chair
Michael Hatzopoulos	Proceedings Chair
Klemens Böhm	Panel and Tutorial Chair
Alfons Kemper	Industrial and Applications Chair
Torsten Grust	Workshop Chair
Christian Böhm	Demonstrations Chair

Table of Contents

Invited Lectures

Database Architecture Fertilizers: Just-in-Time, Just-Enough, and Autonomous Growth <i>Martin Kersten</i>	1
Digital Video: Just Another Data Stream? <i>Alan F. Smeaton</i>	2
Charting a Dataspace: Lessons from Lewis and Clark <i>David Maier</i>	3

Data Streams

Fast Approximate Wavelet Tracking on Streams <i>Graham Cormode, Minos Garofalakis, Dimitris Sacharidis</i>	4
Resource Adaptive Periodicity Estimation of Streaming Data <i>Michail Vlachos, Deepak S. Turaga, Philip S. Yu</i>	23
On Futuristic Query Processing in Data Streams <i>Charu C. Aggarwal</i>	41

Semantic Heterogeneity

Detecting Similarities in Ontologies with the SOQA-SimPack Toolkit <i>Patrick Ziegler, Christoph Kiefer, Christoph Sturm, Klaus R. Dittrich, Abraham Bernstein</i>	59
Holistic Schema Matching for Web Query Interfaces <i>Weifeng Su, Jiyang Wang, Frederick Lochovsky</i>	77
Data Mapping as Search <i>George H.L. Fletcher, Catharine M. Wyss</i>	95

Distributed Databases

Parallelizing Skyline Queries for Scalable Distribution <i>Ping Wu, Caijie Zhang, Ying Feng, Ben Y. Zhao, Divyakant Agrawal, Amr El Abbadi</i>	112
Replication, Load Balancing and Efficient Range Query Processing in DHTs <i>Theoni Pitoura, Nikos Ntarmos, Peter Triantafillou</i>	131
IQN Routing: Integrating Quality and Novelty in P2P Querying and Ranking <i>Sebastian Michel, Matthias Bender, Peter Triantafillou, Gerhard Weikum</i>	149

Multidimensionality and Nearest-Neighbor Searches

Efficient Quantile Retrieval on Multi-dimensional Data <i>Man Lung Yiu, Nikos Mamoulis, Yufei Tao</i>	167
Fast Nearest Neighbor Search on Road Networks <i>Haibo Hu, Dik Lun Lee, Jianliang Xu</i>	186
Approximation Techniques to Enable Dimensionality Reduction for Voronoi-Based Nearest Neighbor Search <i>Christoph Brochhaus, Marc Wichterich, Thomas Seidl</i>	204

Privacy and Security

Authorization-Transparent Access Control for XML Under the Non-Truman Model <i>Yaron Kanza, Alberto O. Mendelzon, Renée J. Miller, Zheng Zhang</i>	222
On Honesty in Sovereign Information Sharing <i>Rakesh Agrawal, Evimaria Terzi</i>	240

Temporal Data Management

Multi-dimensional Aggregation for Temporal Data <i>Michael Böhlen, Johann Gamper, Christian S. Jensen</i>	257
------------------------------------------------------------------------------------------------------------------------	-----

Similarity Search on Time Series Based on Threshold Queries <i>Johannes Aßfalg, Hans-Peter Kriegel, Peer Kröger, Peter Kunath, Alexey Pryakhin, Matthias Renz</i>	276
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Supporting Temporal Slicing in XML Databases <i>Federica Mandreoli, Riccardo Martoglia, Enrico Ronchetti</i>	295
-----------------------------------------------------------------------------------------------------------------------	-----

Text Databases and Information Retrieval

Indexing Shared Content in Information Retrieval Systems <i>Andrei Z. Broder, Nadav Eiron, Marcus Fontoura, Michael Herscovici, Ronny Lempel, John McPherson, Rumping Qi, Eugene Shekita</i>	313
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Feedback-Driven Structural Query Expansion for Ranked Retrieval of XML Data <i>Ralf Schenkel, Martin Theobald</i>	331
----------------------------------------------------------------------------------------------------------------------------	-----

Expressiveness and Performance of Full-Text Search Languages <i>Chavdar Botev, Sihem Amer-Yahia, Jayavel Shanmugasundaram</i>	349
--------------------------------------------------------------------------------------------------------------------------------------------	-----

Schema Management

Model-Independent Schema and Data Translation <i>Paolo Atzeni, Paolo Cappellari, Philip A. Bernstein</i>	368
-------------------------------------------------------------------------------------------------------------------	-----

Physical Design Refinement: The “Merge-Reduce” Approach <i>Nicolas Bruno, Surajit Chaudhuri</i>	386
----------------------------------------------------------------------------------------------------------	-----

Online, Non-blocking Relational Schema Changes <i>Jørgen Løland, Svein-Olaf Hvasshovd</i>	405
----------------------------------------------------------------------------------------------------	-----

Approximation and Estimation

Deferred Maintenance of Disk-Based Random Samples <i>Rainer Gemulla, Wolfgang Lehner</i>	423
---------------------------------------------------------------------------------------------------	-----

Exploiting Cluster Analysis for Constructing Multi-dimensional Histograms on Both Static and Evolving Data <i>Filippo Furfaro, Giuseppe M. Mazzeo, Cristina Sirangelo</i>	442
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

HASE: A Hybrid Approach to Selectivity Estimation for Conjunctive Predicates
Xiaohui Yu, Nick Koudas, Calisto Zuzarte 460

Data and Query Patterns

On High Dimensional Skylines
Chee-Yong Chan, H.V. Jagadish, Kian-Lee Tan, Anthony K.H. Tung, Zhenjie Zhang 478

From Analysis to Interactive Exploration: Building Visual Hierarchies from OLAP Cubes
Svetlana Vinnik, Florian Mansmann 496

DPTree: A Distributed Pattern Tree Index for Partial-Match Queries in Peer-to-Peer Networks
Dyce Jing Zhao, Dik Lun Lee, Qiong Luo 515

XML Queries and Updates

A Decomposition-Based Probabilistic Framework for Estimating the Selectivity of XML Twig Queries
Chao Wang, Srinivasan Parthasarathy, Ruoming Jin 533

Conflicting XML Updates
Mukund Raghavachari, Oded Shmueli 552

Improving the Efficiency of XPath Execution on Relational Systems
Haris Georgiadis, Vasilis Vassalos 570

Data Streams and Pub/Sub Systems

Bridging Physical and Virtual Worlds: Complex Event Processing for RFID Data Streams
Fusheng Wang, Shaorong Liu, Peiya Liu, Yijian Bai 588

On Concurrency Control in Sliding Window Queries over Data Streams
Lukasz Golab, Kumar Gaurav Bijay, M. Tamer Özsu 608

Towards Expressive Publish/Subscribe Systems
Alan Demers, Johannes Gehrke, Mingsheng Hong, Mirek Riedewald, Walker White 627

Data Mining and Knowledge Discovery

Finding Data Broadness Via Generalized Nearest Neighbors <i>Jayendra Venkateswaran, Tamer Kahveci, Orhan Camoglu</i>	645
TrajPattern: Mining Sequential Patterns from Imprecise Trajectories of Mobile Objects <i>Jiong Yang, Meng Hu</i>	664
On Exploring the Power-Law Relationship in the Itemset Support Distribution <i>Kun-Ta Chuang, Jiun-Long Huang, Ming-Syan Chen</i>	682

Images, Multimedia, and User Interfaces

Fast Query Point Movement Techniques with Relevance Feedback for Content-Based Image Retrieval <i>Danzhou Liu, Kien A. Hua, Khanh Vu, Ning Yu</i>	700
On Fast Non-metric Similarity Search by Metric Access Methods <i>Tomáš Skopal</i>	718
Constructing a Generic Natural Language Interface for an XML Database <i>Yunyao Li, Huahai Yang, H.V. Jagadish</i>	737

XML Data Management

A New Design for a Native XML Storage and Indexing Manager <i>Jihad Boulos, Shant Karakashian</i>	755
XML Duplicate Detection Using Sorted Neighborhoods <i>Sven Puhlmann, Melanie Weis, Felix Naumann</i>	773
Handling Interlinked XML Instances on the Web <i>Erik Behrends, Oliver Fritzen, Wolfgang May</i>	792

Query Optimization

Query Planning in the Presence of Overlapping Sources <i>Jens Bleiholder, Samir Khuller, Felix Naumann, Louiqa Raschid, Yao Wu</i>	811
-------------------------------------------------------------------------------------------------------------------------------------------------	-----

Optimizing Monitoring Queries over Distributed Data
Frank Neven, Dieter Van de Craen 829

Progressive Query Optimization for Federated Queries
Stephan Ewen, Holger Kache, Volker Markl, Vijayshankar Raman 847

Data Structures and Indexing

Indexing Spatially Sensitive Distance Measures Using Multi-resolution Lower Bounds
Vebjorn Ljosa, Arnab Bhattacharya, Ambuj K. Singh 865

Indexing Incomplete Databases
Guadalupe Canahuate, Michael Gibas, Hakan Ferhatosmanoglu 884

FlexInd: A Flexible and Parameterizable Air-Indexing Scheme for Data Broadcast Systems
André Seifert, Jen-Jou Hung 902

Nontraditional Query Processing

Multi-query SQL Progress Indicators
Gang Luo, Jeffrey F. Naughton, Philip S. Yu 921

Finding Equivalent Rewritings in the Presence of Arithmetic Comparisons
Foto Afrati, Rada Chirkova, Manolis Gergatsoulis, Vassia Pavlaki 942

Fast Computation of Reachability Labeling for Large Graphs
Jiefeng Cheng, Jeffrey Xu Yu, Xuemin Lin, Haixun Wang, Philip S. Yu 961

Spatial Data Management

Distributed Spatial Clustering in Sensor Networks
Anand Meka, Ambuj K. Singh 980

SCUBA: Scalable Cluster-Based Algorithm for Evaluating Continuous Spatio-temporal Queries on Moving Objects
Rimma V. Nehme, Elke A. Rundensteiner 1001

Caching Complementary Space for Location-Based Services <i>Ken C.K. Lee, Wang-Chien Lee, Baihua Zheng, Jianliang Xu</i>	1020
--------------------------------------------------------------------------------------------------------------------------------------	------

Extending Data Base Technology

Evolving Triggers for Dynamic Environments <i>Goce Trajcevski, Peter Scheuermann, Oliviu Ghica, Annika Hinze, Agnes Voisard</i>	1039
A Framework for Distributed XML Data Management <i>Serge Abiteboul, Ioana Manolescu, Emanuel Taropa</i>	1049
Querying and Updating Probabilistic Information in XML <i>Serge Abiteboul, Pierre Senellart</i>	1059
An ECA Rule Rewriting Mechanism for Peer Data Management Systems <i>Dan Zhao, John Mylopoulos, Iluju Kiringa, Verena Kantere</i>	1069

Industrial Session

Business Intelligence

A Metric Definition, Computation, and Reporting Model for Business Operation Analysis <i>Fabio Casati, Malu Castellanos, Umeshwar Dayal, Ming-Chien Shan</i>	1079
BISON: Providing Business Information Analysis as a Service <i>Hakan Hacigümüş, James Rhodes, Scott Spangler, Jeffrey Kreulen</i>	1084
The Design and Architecture of the τ -Synopses System <i>Yossi Matias, Leon Portman, Natasha Drukh</i>	1088

Database System Enhancements

Integrating a Maximum-Entropy Cardinality Estimator into DB2 UDB <i>Marcel Kutsch, Peter J. Haas, Volker Markl, Nimrod Megiddo, Tam Minh Tran</i>	1092
----------------------------------------------------------------------------------------------------------------------------------------------------------------	------

Improving DB2 Performance Expert - A Generic Analysis Framework
Laurent Mignet, Jayanta Basak, Manish Bhide, Prasan Roy, Sourashis Roy, Vibhuti S. Sengar, Ranga R. Vatsavai, Michael Reichert, Torsten Steinbach, D.V.S. Ravikant, Soujanya Vadapalli 1097

Managing Collections of XML Schemas in Microsoft SQL Server 2005
Shankar Pal, Dragan Tomic, Brandon Berg, Joe Xavier 1102

Business Data Processing

Enabling Outsourced Service Providers to Think Globally While Acting Locally
Kevin Wilkinson, Harumi Kuno, Kannan Govindarajan, Kei Yuasa, Kevin Smathers, Jyotirmaya Nanda, Umeshwar Dayal 1106

Another Example of a Data Warehouse System Based on Transposed Files
Antonio Albano, Luca De Rosa, Lucio Goglia, Roberto Goglia, Vincenzo Minei, Cristian Dumitrescu 1110

Demonstration Papers

XG: A Grid-Enabled Query Processing Engine
Radu Sion, Ramesh Natarajan, Inderpal Narang, Thomas Phan 1115

Managing and Querying Versions of Multiversion Data Warehouse
Robert Wrembel, Tadeusz Morzy 1121

Natix Visual Interfaces
A. Böhm, M. Brantner, C-C. Kanne, N. May, G. Moerkotte 1125

Hermes - A Framework for Location-Based Data Management
Nikos Pelekis, Yannis Theodoridis, Spyros Vosinakis, Themis Panayiotopoulos 1130

TeNDaX, a Collaborative Database-Based Real-Time Editor System
Stefania Leone, Thomas B. Hodel-Widmer, Michael Boehlen, Klaus R. Dittrich 1135

Synopses Reconciliation Via Calibration in the τ -Synopses System <i>Yariv Matia, Yossi Matias, Leon Portman</i>	1139
X-Evolution: A System for XML Schema Evolution and Document Adaptation <i>Marco Mesiti, Roberto Celle, Matteo A. Sorrenti, Giovanna Guerrini</i>	1143
TQuEST: Threshold Query Execution for Large Sets of Time Series <i>Johannes Aßfalg, Hans-Peter Kriegel, Peer Kröger, Peter Kunath, Alexey Pryakhin, Matthias Renz</i>	1147
VICO: Visualizing Connected Object Orderings <i>Stefan Brecheisen, Hans-Peter Kriegel, Matthias Schubert, Michael Gruber</i>	1151
XQueryViz: An XQuery Visualization Tool <i>Jihad Boulos, Marcel Karam, Zeina Koteiche, Hala Ollaic</i>	1155
SAT: Spatial Awareness from Textual Input <i>Dmitri V. Kalashnikov, Yiming Ma, Sharad Mehrotra, Ramaswamy Hariharan, Nalini Venkatasubramanian, Naveen Ashish</i>	1159
MUSCLE: Music Classification Engine with User Feedback <i>Stefan Brecheisen, Hans-Peter Kriegel, Peter Kunath, Alexey Pryakhin, Florian Vorberger</i>	1164
iMONDRIAN: A Visual Tool to Annotate and Query Scientific Databases <i>Floris Geerts, Anastasios Kementsietsidis, Diego Milano</i>	1168
The SIRUP Ontology Query API in Action <i>Patrick Ziegler, Christoph Sturm, Klaus R. Dittrich</i>	1172
Querying Mediated Geographic Data Sources <i>Mehdi Essid, François-Marie Colonna, Omar Boucelma, Abdelkader Betari</i>	1176
FIS-by-Step: Visualization of the Fast Index Scan for Nearest Neighbor Queries <i>Elke Achttert, Dominik Schwald</i>	1182

ArHeX: An Approximate Retrieval System for Highly Heterogeneous XML Document Collections
Ismael Sanz, Marco Mesiti, Giovanna Guerrini, Rafael Berlanga Llavori 1186

MonetDB/XQuery—Consistent and Efficient Updates on the Pre/Post Plane
Peter Boncz, Jan Flokstra, Torsten Grust, Maurice van Keulen, Stefan Manegold, Sjoerd Mullender, Jan Rittinger, Jens Teubner 1190

STRIDER: A Versatile System for Structural Disambiguation
Federica Mandreoli, Riccardo Martoglia, Enrico Ronchetti 1194

An Extensible, Distributed Simulation Environment for Peer Data Management Systems
Katja Hose, Andreas Job, Marcel Karnstedt, Kai-Uwe Sattler 1198

Panel

Data Management in the Social Web
Karl Aberer 1203

Author Index 1205