

# Advances in Database Technology – EDBT 2009

---

12th International Conference  
on Extending Database Technology  
Saint Petersburg, March 24–26, 2009  
Proceedings

Editors:

Martin Kersten (CWI, The Netherlands)  
Boris Novikov (University of Saint Petersburg, Russia)  
Jens Teubner (ETH Zurich, Switzerland)  
Vladimir Polutin (HP Labs, Russia)  
Stefan Manegold (CWI, The Netherlands)

Advances in Database Technology – EDBT 2009  
Proceedings of the 12th International Conference  
on Extending Database Technology  
Saint Petersburg, Russia, March 24–26, 2009

Editors:  
Martin Kersten  
Boris Novikov  
Jens Teubner  
Vladimir Polutin  
Stefan Manegold

The Association for Computing Machinery  
2 Penn Plaza, Suite 701  
New York, NY, 10121-0701

ACM COPYRIGHT NOTICE. Copyright © 2009 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, +1-978-750-8400, +1-978-750-4470 (fax).

ACM ISBN: 978-1-60558-422-5

# Table of Contents

Foreword .....	xi
Program Committee Members .....	xii–xiii

## Invited Papers

Data Integration Flows for Business Intelligence <i>Umeshwar Dayal, Malu Castellanos, Alkis Simitsis, and Kevin Wilkinson</i> .....	1–11
--	------

## Research Sessions

### System Architectures

Optimized Union of Non-disjoint Distributed Data Sets <i>Itay Dar, Tova Milo, and Elad Verbin</i> .....	12–23
Shore-MT: A Scalable Storage Manager for the Multicore Era <i>Ryan Johnson, Ippokratis Pandis, Nikos Hardavellas, Anastasia Ailamaki, and Babak Falsafi</i> .....	24–35
Workload-Aware Data Partitioning in Community-Driven Data Grids <i>Tobias Scholl, Bernhard K. Bauer, Jessica Müller, Benjamin Gufler, Angelika Reiser, and Alfons Kemper</i> .....	36–47

### Spatio-Temporal

Sequenced, Spatio-Temporal Aggregation in Road Networks <i>Igor Timko, Michael Boehlen, and Johann Gamper</i> .....	48–59
Processing Probabilistic Spatio-Temporal Range Queries over Moving Objects with Uncertainty <i>Bruce Chung, Wang-Chien Lee, and Arbee L.P. Chen</i> .....	60–71
Anonymizing Moving Objects: How to Hide a MOB in a Crowd? <i>Roman Yarovoy, Francesco Bonchi, Laks V.S. Lakshmanan, and Hui Wang</i> .....	72–83

### Database Summarization

Type-Based Categorization of Relational Attributes <i>Babak Ahmadi, Marios Hadjieleftheriou, Thomas Seidl, Divesh Srivastava, and Suresh Venkatasubramanian</i> .....	84–95
AlphaSum: Size-Constrained Table Summarization using Value Lattices <i>K. Selcuk Candan, Huiping Cao, Yan Qi, and Maria Luisa Sapino</i> .....	96–107
Answering Aggregate Keyword Queries on Relational Databases Using Minimal Group-bys <i>Bin Zhou and Jian Pei</i> .....	108–119

## Query Processing

### Rule-Based Multi-Query Optimization

*Mingsheng Hong, Mirek Riedewald, Christoph Koch, Johannes Gehrke, and Alan Demers* ..... 120–131

### Managing Long-Running Queries

*Stefan Krompaß, Harumi Kuno, Janet Wiener, Kevin Wilkinson, Umeshwar Dayal, and Alfons Kemper* ..... 132–143

### Continuous Visible Nearest Neighbor Queries

*Yunjun Gao, Baihua Zheng, Wang-Chien Lee, and Gencai Chen* ..... 144–155

## XML, XPath, XQuery

### Query Ranking in Probabilistic XML Data

*Lijun Chang, Jeffrey Xu Yu, and Lu Qin* ..... 156–167

### On Rewriting XPath Queries Using Views

*Foto Afrati, Rada Chirkova, Manolis Gergatsoulis, Vassia Pavlaki, Benny Kimelfeld, and Yehoshua Sagiv* ..... 168–179

### Parallelization of XPath Queries using Multi-core Processors: Challenges and Experiences

*Lipyeow Lim, Rajesh Bordawekar, and Oded Shmueli* ..... 180–191

## Graph Techniques

### GADDI: Distance Index based Subgraph Matching in Biological Networks

*Shijie Zhang, Shirong Li, and Jiong Yang* ..... 192–203

### A Novel Approach for Efficient Supergraph Query Processing on Graph Databases

*Shuo Zhang, Jianzhong Li, Hong Gao, and Zhaonian Zou* ..... 204–215

### Flexible Query Answering on Graph-modeled Data

*Giorgio Villani, Federica Mandreoli, Riccardo Martoglia, and Wilma Penzo* ..... 216–227

## Privacy & Security

### Privacy-Preserving Data Mashup

*Noman Mohammed, Benjamin C. M. Fung, Ke Wang, and Patrick C. K. Hung* .. 228–239

### On the Comparison of Microdata Disclosure Control Algorithms

*Rinku Dewri, Indrajit Ray, Indrakshi Ray, and Darrell Whitley* ..... 240–251

### Detecting Privacy Violations in Database Publishing using Disjoint Queries

*Millist Vincent, Mukesh Mohania, and Mizuho Iwaihara* ..... 252–262

## Data Models

### On Keys, Foreign Keys and Nullable Attributes in Relational Mapping Systems

*Luca Cabibbo* ..... 263–274

### A runtime approach to model-independent schema and data translation

*Paolo Atzeni, Luigi Bellomarini, Francesca Bugiotti, and Giorgio Gianforme* ..... 275–286

### A methodology for preference-based personalization of contextual data

*Antonio Miele, Elisa Quintarelli, and Letizia Tanca* ..... 287–298

## Stream Processing

- Scalable Stream Join Processing with Expensive Predicates: Workload Distribution and Adaptation by Time-Slicing  
*Song Wang and Elke Rundensteiner* ..... 299–310
- Indexing Density Models for Incremental Learning and Anytime Classification on Data Streams  
*Thomas Seidl, Ira Assent, Philipp Kranen, Ralph Krieger, and Jennifer Herrmann* 311–322
- Exploiting the Power of Relational Databases for Efficient Stream Processing  
*Erietta Liarou, Romulo Goncalves, and Stratos Idreos* ..... 323–334

## XML, XPath, XQuery

- A Sampling Approach for XML Query Selectivity Estimation  
*Cheng Luo, Zhewei Jiang, Wen-Chi Hou, Feng Yu, and Qiang Zhu* ..... 335–344
- Recursion in XQuery: Put Your Distributivity Safety Belt On  
*Loredana Afanasiev, Torsten Grust, Maarten Marx, Jan Rittinger, and Jens Teubner* ..... 345–356
- Expressive, yet Tractable XML Keys  
*Sven Hartmann and Sebastian Link* ..... 357–367

## Database Summarization

- It Takes Variety to Make a World: Diversification in Recommender Systems  
*Cong Yu, Laks V.S. Lakshmanan, and Sihem Amer-Yahia* ..... 368–378
- Supporting Annotations on Relations  
*Mohamed Eltabakh, Walid G. Aref, Ahmed Elmagarmid, Mourad Ouzzani, and Yasin Silva* ..... 379–390
- DataClouds: Summarizing Keyword Search Results over Structured Data  
*Georgia Koutrika, Zahra Mohammadi Zadeh, and Hector Garcia-Molina* ..... 391–402

## Query Processing

- Sample Synopses for Approximate Answering of Group-By Queries  
*Philipp Rösch and Wolfgang Lehner* ..... 403–414
- A Query Processor for Prediction-Based Monitoring of Data Streams  
*Sergio Ilarri, Ouri Wolfson, Eduardo Mena, Arantza Illarramendi, and Prasad Sistla* ..... 415–426
- Flower-CDN: A Hybrid P2P Overlay for Efficient Query Processing in CDN  
*Manal El Dick, Esther Pacitti, and Bettina Kemme* ..... 427–438

## Top-K Techniques

- Zerber+R: Top-k Retrieval from a Confidential Index  
*Sergej Zerr, Daniel Olmedilla, Wolfgang Nejdl, and Wolf Siberski* ..... 439–449
- Efficient Top-K count queries over imprecise duplicates  
*Sunita Sarawagi, Vinay Deshpande, and Sourabh Kasliwal* ..... 450–461
- The C-ND Tree: A Multidimensional Index for Hybrid Continuous and Non-ordered Discrete Data Spaces  
*Changqing Chen, Sakti Pramanik, Qiang Zhu, Alok Watve, and Gang Qian* ..... 462–471

## Graph Techniques

G-Hash: Towards Fast Kernel-based Similarity Search in Large Graph Databases <i>Xiaohong Wang, Aaron Smalter, Jun Huan, and Gerald Lushington</i> .....	472–480
On-line Exact Shortest Distance Query Processing <i>Cheng Jiefeng and Jeffrey Xu Yu</i> .....	481–492
Efficiently Indexing Shortest Paths by Exploiting Symmetry in Graphs <i>Yanghua Xiao, Wentao Wu, Jian Pei, Wei Wang, and Zhenying He</i> .....	493–504

## Data Mining

Estimating the Number of Frequent Itemsets in a Large Database <i>Ruoming Jin, Scott McCallen, Yuri Breitbart, David Fuhry, and Dong Wang</i> .....	505–516
FOGGER: An Algorithm for Graph Generator Discovery <i>Zhiping Zeng, Jianyong Wang, Jun Zhang, and Lizhu Zhou</i> .....	517–528
Neighbor-Based Pattern Detection for Windows Over Streaming Data <i>Di Yang, Elke Rundensteiner, and Matthew Ward</i> .....	529–540
Efficient Constraint Evaluation in Categorical Sequential Pattern Mining for Trajectory Databases <i>Leticia Gomez and Alejandro Vaisman</i> .....	541–552

## Heterogeneous & Distributed

Flexible and Efficient Querying and Ranking on Hyperlinked Data Sources <i>Ramakrishna Varadarajan, Hector Rodriguez-Drumond, Vagelis Hristidis, Louiqa Raschid, Maria-Esther Vidal, and Luis Daniel Ibáñez</i> .....	553–564
RankClus: Integrating Clustering with Ranking for Heterogenous Information Network Analysis <i>Yizhou Sun, Jiawei Han, Peixiang Zhao, Zhijun Yin, Hong Cheng, and Tianyi Wu</i> .....	565–576
Evaluating Very Large Datalog Queries on Social Networks <i>Royi Ronen and Oded Shmueli</i> .....	577–587

## System Architectures

A Sequential Indexing Scheme for Flash-Based Embedded Systems <i>Shaoyi Yin, Philippe Pucheral, and Xiaofeng Meng</i> .....	588–599
Secondary Bitmap Indexes with Vertical and Horizontal Partitioning <i>Guadalupe Canahuate, Tan Apaydin, Ahmet Sacan, and Hakan Ferhatosmanoglu</i> .	600–611
Automating the Loading of Business Process Data Warehouses <i>Malu Castellanos, Alkis Simitsis, Kevin Wilkinson, and Umeshwar Dayal</i> .....	612–623

## Privacy & Security

Hiding Distinguished Ones into Crowd: Privacy-Preserving Publishing Data with Outliers <i>Hui Wang and Ruilin Liu</i> .....	624–635
An Efficient Online Auditing Approach to Limit Private Data Disclosure <i>Haibing Lu, Yingjiu Li, Vijayalakshmi Athuri, and Jaideep Vaidya</i> .....	636–647
Continuous Privacy Preserving Publishing of Data Streams <i>Bin Zhou, Yi Han, Jian Pei, Bin Jiang, Yufei Tao, and Yan Jia</i> .....	648–659

## Uncertainty

Top-k Dominating Queries in Uncertain Databases <i>Xiang Lian and Lei Chen</i> .....	660–671
Evaluating Probability Threshold k-Nearest-Neighbor Queries over Uncertain Data <i>Reynold Cheng, Lei Chen, Jinchuan Chen, and Xike Xie</i> .....	672–683
PROUD: A Probabilistic Approach to Processing Similarity Queries over Uncertain Data Streams <i>MI-YEN YEH, Philip Yu, Kun-Lung Wu, and Ming-Syan Chen</i> .....	684–695

## Workflow Techniques

Fair, Effective, Efficient and Differentiated Scheduling in an Enterprise Data Warehouse <i>Chetan Gupta, Abhay Mehta, Song Wang, and Umeshwar Dayal</i> .....	696–707
Efficient Identification of Starters and Followers in Social Media <i>Michael Mathioudakis and Nick Koudas</i> .....	708–719
A Data Damage Tracking Quarantine and Recovery (DTQR) Scheme for Mission-Critical Database Systems <i>Kun Bai and Peng Liu</i> .....	720–731

## Multi-Dimensional

Unrestricted Wavelet Synopses under Maximum Error Bound <i>Chaoyi Pang, Qing Zhang, David Hansen, and Anthony Maeder</i> .....	732–743
Distributed Similarity Search in High Dimensions Using Locality Sensitive Hashing <i>Parisa Haghani, Sebastian Michel, and Karl Aberer</i> .....	744–755
Multiplicative Synopses for Relative-Error Metrics <i>Panagiotis Karras</i> .....	756–767
LCS-Hist: Taming Massive High-Dimensional Data Cube Compression <i>Alfredo Cuzzocrea and Paolo Serafino</i> .....	768–779

## Caching Techniques

Caching Content-based Queries for Robust and Efficient Image Retrieval <i>Fabrizio Falchi, Claudio Lucchese, Salvatore Orlando, Raffaele Perego, and Fausto Rabitti</i> .....	780–790
An Approach for Detecting Relevant Updates to Cached Data Using XML and Active Databases <i>Essam Mansour and Hagen Höpfner</i> .....	791–802
Self-Tuning Query Mesh for Adaptive Multi-Route Query Processing <i>Rimma Nehme, Elke Rundensteiner, and Elisa Bertino</i> .....	803–814

## Information Retrieval

Retrieving Meaningful Relaxed Tightest Fragments for XML Keyword Search <i>LINGBO KONG, Rémi Gilleron, and Aurélien Lemay</i> .....	815–826
Approximate Substring Selectivity Estimation <i>Hongrae Lee, Raymond Ng, and Kyuseok Shim</i> .....	827–838
Finding Frequent Co-occurring Terms in Relational Keyword Search <i>Yufei Tao and Jeffrey Xu Yu</i> .....	839–850

## Query Processing

Time-completeness trade-offs in record linkage using Adaptive Query Processing <i>Roald Lengu, Paolo Missier, Alvaro Fernandes, Giovanna Guerrini, and Marco Mesiti</i> .....	851–861
Interactive Query Refinement <i>Chaitanya Mishra and Nick Koudas</i> .....	862–873
Continuous Probabilistic Nearest-Neighbor Queries for Uncertain Trajectories <i>Goce Trajcevski, Roberto Tamassia, Hui Ding, Peter Scheuermann, and Isabel Cruz</i> .....	874–885

## Top-K Techniques

Reverse k-Nearest Neighbor Search in Dynamic and General Metric Databases <i>Elke Achtert, Hans-Peter Kriegel, Peer Kröger, Matthias Renz, and Andreas Züfle</i> .....	886–897
Top-k Dominant Web Services Under Multi-Criteria Matching <i>Dimitrios Skoutas, Dimitris Sacharidis, Alkis Simitsis, Verena Kantere, and Timos Sellis</i> .....	898–909
Ranking Objects Based on Relationships and Fixed Associations <i>Albert Angel, Surajit Chaudhuri, Gautam Das, and Nick Koudas</i> .....	910–921

## Potpourri

Towards Integrated and Efficient Scientific Sensor Data Processing: A Database Approach <i>Ji Wu, Yongluan Zhou, Karl Aberer, and Kian-Lee Tan</i> .....	922–933
Flexible and Scalable Storage Management for Data-intensive Stream Processing <i>Irina Botan, Gustavo Alonso, Peter Fischer, Donald Kossmann, and Nesime Tatbul</i> .....	934–945
A view selection algorithm with performance guarantee <i>Sofian Maabout, Nicolas Hanusse, and Radu Tofan</i> .....	946–957

## Provenance

Efficient Provenance Storage over Nested Data Collections <i>Manish Anand, Shawn Bowers, Timothy McPhillips, and Bertram Ludaescher</i> ....	958–969
Schema-Conscious Filtering of XML Documents <i>Panu Silvast, Seppo Sippu, and Eljas Soisalon-Soininen</i> .....	970–981
Provenance for nested subqueries <i>Boris Glavic and Gustavo Alonso</i> .....	982–993

## Spatio-Temporal

A Data Model For Trip Planning in Multimodal Transportation Systems <i>Joel Booth, Ouri Wolfson, Isabel Cruz, and Prasad Sistla</i> .....	994–1005
Parsimonious Temporal Aggregation <i>Juozas Gordevicius, Johann Gamper, and Michael Boehlen</i> .....	1006–1017
Fast Object Search on Road Networks <i>Ken C.K. Lee, Wang-Chien Lee, and Baihua Zheng</i> .....	1018–1029



## Skylines

Finding the Influence Set through Skylines <i>Xiaobing Wu, Yufei Tao, Raymond Chi-Wing Wong, Ling Ding, and Jeffrey Xu Yu</i>	1030–1041
Efficient Skyline Computation in Metric Space <i>David Fuhry, Ruoming Jin, and Donghui Zhang</i>	1042–1051
Efficient Skyline Retrieval with Arbitrary Similarity Measures <i>Deepak P, Prasad Deshpande, Debapriyo Majumdar, and Raghu Krishnapuram</i>	1052–1063
<b>Transaction Processing</b>	
Transactions on the Multiversion B-Tree <i>Tuukka Haapasalo, Ibrahim Jaluta, Bernhard Seeger, Seppo Sippu, and Eljas Soisalon-Soininen</i>	1064–1075
Efficient maintenance techniques for views over active documents <i>Serge Abiteboul, Pierre Bourhis, and Bogdan Marinoiu</i>	1076–1087
Towards Materialized View Selection for Distributed Databases <i>Leonardo Weiss Ferreira Chaves, Erik Buchmann, Fabian Hueske, and Klemens Böhm</i>	1088–1099

## Industrial Sessions

### Industrial Session

Personalizing Entity Detection and Recommendation with a Fusion of Web Log Mining Techniques <i>Kathleen Tsoukalas, Bin Zhou, Jian Pei, and Davor Cubranic</i>	1100–1103
Estimating Aggregates in Time-Constrained Approximate Queries in Oracle <i>Ying Hu, Seema Sundara, and Jagannathan Srinivasan</i>	1104–1107

## Demonstrations

### Demo Group 1

BaseX & DeepFS - Joint Storage for Filesystem and Database <i>Alexander Holupirek, Christian Grün, and Marc H. Scholl</i>	1108–1111
Xoom: A tool for zooming in and out of XML documents <i>Maya Ramanath and Kondreddi Sarath Kumar</i>	1112–1115
HIDE: Heterogeneous Information DE-identification <i>James Gardner, Li Xiong, Kanwei Li, and James J. Lu</i>	1116–1119
MVT: A Schema Mapping Validation Tool <i>Guillem Rull, Carles Farré, Ernest Teniente, and Toni Urpí</i>	1120–1123
A Tool for Mapping Discovery over Revealing Schemas <i>Verena Kantere, Dimos Bousounis, and Timos Sellis</i>	1124–1127
GCIP: Exploiting the Generation and Optimization of Integration Processes <i>Matthias Boehm, Dirk Habich, Wolfgang Lehner, and Uwe Wloka</i>	1128–1131

## Demo Group 2

CourseCloud: Summarizing and Refining Keyword Searches over Structured Data <i>Georgia Koutrika, Zahra Mohammadi Zadeh, and Hector Garcia-Molina</i> .....	1132–1135
EventSummarizer: A tool for summarizing large event sequences <i>Jerry Kiernan and Evimaria Terzi</i> .....	1136–1139
High-Performance Information Extraction with AliBaba <i>Peter Palaga, Long Nguyen, Ulf Leser, and Jörg Hakenberg</i> .....	1140–1143
Exploiting Similarity-aware Grouping in Decision Support Systems <i>Yasin Silva, Muhammad Arshad, and Walid G. Aref</i> .....	1144–1147
MarcoPolo: A Community System for Sharing and Integrating Travel Information on Maps <i>Yueguo Chen, Su Chen, Yu Gu, Mei Hui, Feng Li, Chen Liu, Liangxu Liu, Beng Chin Ooi, Xiaoyan Yang, Dongxiang Zhang, and Yuan Zhou</i> .....	1148–1151
NNexus: An Automatic Linker for Collaborative Web-Based Corpora <i>James Gardner, Aaron Krowne, and Li Xiong</i> .....	1152–1155

## Tutorials

Performance Evaluation in Database Research: Principles and Experience <i>Stefan Manegold and Ioana Manolescu</i> .....	1156
Geographic Privacy-aware Knowledge Discovery and Delivery <i>Fosca Giannotti, Dino Pedreschi, and Yannis Theodoridis</i> .....	1157–1158
Scalable OLAP and Mining of Information Networks <i>Jiawei Han, Xifeng Yan, and Philip Yu</i> .....	1159
Author Index .....	1160–1163

# Foreword

EDBT 2009 marks a turning point in the way this conference is ran. It is the first one that breaks with the two year cycle, offering a yearly platform for our community. Aside from teaming up with the ICDT, there has been a few major breakthroughs: number of submissions, review process innovations, and program size.

**Submissions.** We received 408 abstracts for the main track, 32 demo papers, and 7 short industrial papers. Numbers to be proud off, as they align with other international conferences in our field. The deadlines for abstract submission, final paper upload and notification were aligned with the neighboring conferences ICDE and ACM SIGMOD.

The industrial papers have been purposely limited to four pages. Research results from the industrial labs are on par with those of the research institutes. They don't need special treatment and this enables the industrial track to be used for what it originally meant, to provide a stage for product innovations. This year the panel sessions were dropped to favor acceptance of more good research papers. The workshops around the conference are often a well-targeted platform to explore new areas of research.

A common practice in our field is to roll rejected papers forward to the next conference. This leads to an ever growing pool of good papers or lottery submissions. We considered reducing the review load somewhat by also rolling forward paper assignments from ICDE. Due to time and logistic constraints we could not implement this strategy, but recommend it to future conferences. It would reduce the paper pool and increase the quality.

**Review process.** The review process was organised using the Conference Management Toolkit (CMT) sponsored by Microsoft Research. It has become a standard system for handling such large number of papers and reviews. The review process was purposely set tight with only six weeks for the reviewers to read, contemplate the contributions, and upload their reports. In line with previous years, the bulk (80%) of these were received around the review deadline.

Another innovation was to give advanced notice during the discussion phase to authors whose paper most likely would not make it into the program. It was considered a gesture to these unfortunate authors and gave them an extra week to work on improvements for the next conference. Unfortunately, CMT did not allow also selective disclosure of the review reports yet.

**Program structure.** In line with recent conferences the pool of high-quality papers has grown, pushing the conference organisers to maximise the allotment of space to provide a stage for our research.

Saint Petersburg is a historical city where you can spent days in the Hermitage. This grandeur is extended to the location offered for the conference. It allowed for four parallel streams! A total of 2 keynotes, 92 high-quality papers, three tutorials, one industrial session and 12 demos could be accommodated in a tightly packed three day meeting. It marks EDBT 2009 as a memorable year of great scientific harvesting. A side visit to ICDT or extended stay to visit the workshops all make this trip scientifically and culturally rewarding.

Martin Kersten (Program Chair)  
Boris Novikov (General Chair)  
Jens Teubner (Proceedings Chair)

Vladimir Polutin (Industrial & Appl. Chair)  
Stefan Manegold (Demonstrations Chair)

# Program Committee Members

## Research

Loredana Afanasiev	Ehud Gudes	Philippe Pucheral
Gustavo Alonso	Ralf Guting	Uwe Rahm
Walid Aref	Hele-Mai Haav	Krithi Ramamritham
Wolf-Tilo Balke	Sven Helmer	Tore Risch
David Bell	Jan Hidders	Peter Scheuermann
Michael Boehlen	Stratos Idreos	Marc H. Scholl
Christian Bohm	H.V. Jagadish	Thomas Seidl
Angela Bonifati	Keith Jeffery	Timos Sellis
Mokrane Bouezeghoub	Ruoming Jin	Arno Siebes
Angelo Brayner	Bettina Kemme	Christian S. Jensen
Francois Bry	Martin Kersten (chair)	Cristina Sirangelo
Alex Buchmann	George Kollios	Nicolas Spiratos
Fabio Casati	Nick Koudas	Dan Suciu
Tiziana Catarci	Georgia Koutrika	Val Tannen
Ugur Cetintemel	Chen Li	Yufei Tao
Zhiyuan Chen	Qing Li	Martin Theobald
Carl-Christian Kanne	Nikos Mamoulis	David Toman
Vassilis Christophidis	Florian Matthes	Hakki Toroslu
Christine Collet	Marta Mattoso	Peter Triantafillou
Sara Comai	Prasenjit Mitra	Can Türker
Umeshwar Dayal	Noureddine Mouaddib	Aparna Varde
Stefan Dessloch	Jose OrlandoPereira	Vasilis Vassalos
Arjen de Vries	Gultekin Ozsoyoglu	Haixun Wang
Jens Dittrich	Tamer Özsu	Kyu-Young Whang
Amr El Abbadi	Esther Pacitti	Ouri Wolfson
Norbert Fuhr	Norman Paton	Jeffrey Xu Yu
Michael Gertz	Jian Pei	Haruo Yokota
Shahram Ghandeharizadeh	Reinhard Pichler	Philip Yu
Bart Goethals	Evaggelia Pitoura	
Torsten Grabs	Sunil Prabhakar	

## Industrial & Applications

Ricardo Baeza-Yates	Lipyew Lim	Ilya Segalovich
Jidong Chen	Pankaj Mehra	Timos Sellis
Dean Jacobs	Vladimir Polutin (chair)	Jens Teubner
Christian Lang	Stefanie Scherzinger	

## Demonstrations

Daniel Abadi	Philippe Bonnet	Ander de Keijzer
Ira Assent	Christof Bornhoevd	Johann Gamper
Bishwaranjan Bhattacharjee	Luc Bouganim	Raghav Kaushik

Qiong Luo  
Stefan Manegold (chair)  
Holger Meyer  
Dan Olteanu

Raghunath Othayoth Nambiar  
Meikel Poess  
Neoklis Polyzotis  
Jun Rao

Ralf Schenkel  
Matthias Schubert  
Alkis Simitsis