

38th International Conference on Very Large Data Bases 2012

(VLDB 2012)

**Proceedings of the VLDB Endowment
Volume 5 Numbers 1-12**

**Istanbul, Turkey
27-31 August 2012**

Volume 1 of 3

Editors:

Z. Meral Özsoyoğlu

**ISBN: 978-1-62276-758-8
ISSN: 2150-8097**

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2012) by the VLDB Endowment
All rights reserved.

Printed by Curran Associates, Inc. (2013)

For permission requests, please contact the VLDB Endowment
at the address below.

VLDB Endowment
Attn: Marianne Siroker
353 Serra Mall – Stanford University
Stanford CA 94305

Phone: (650) 723-0872
Fax: (650) 725-2588

siroker@cs.stanford.edu

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2634
Email: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

Volume 1

NUMBER 1

RESEARCH TRACK PAPERS

Explanation-Based Auditing	1
<i>Daniel Fabbri, Kristen Lefevre</i>	
Human-powered Sorts and Joins	13
<i>Adam Marcus, Eugene Wu, David Karger, Samuel Madden, Robert Miller</i>	
Verifying Computations with Streaming Interactive Proofs	25
<i>Graham Cormode, Justin Thaler, Ke Yi</i>	
A MovingObject Index for Efficient Query Processing with Peer-Wise Location Privacy	37
<i>Dan Lin, Christian S. Jensen, Rui Zhang, Lu Xiao, Jiaheng Lu</i>	
ERA: Efficient Serial and Parallel Suffix Tree Construction for Very Long Strings	49
<i>Essam Mansour, Amin Allam, Spiros Skiadopoulos, Panos Kalnis</i>	
Fast Updates on Read-Optimized Databases Using Multi-Core CPUs	61
<i>Jens Krueger, Changkyu Kim, Martin Grund, Nadathur Satish, David Schwalb, Jatin Chhugani, Hasso Plattner, Pradeep Dubey, Alexander Zeier</i>	
A Data-Based Approach to Social Influence Maximization	73
<i>Amit Goyal, Francesco Bonchi, Laks V. S. Lakshmanan</i>	

NUMBER 2

RESEARCH TRACK PAPERS

On Predictive Modeling for Optimizing Transaction Execution in Parallel OLTP Systems	85
<i>Andrew Pavlo, Evan P. C. Jones, Stanley Zdonik</i>	
View Selection in Semantic Web Databases	97
<i>François Goasdoué, Konstantinos Karanasos, Julien Leblay, Ioana Manolescu</i>	
Building Wavelet Histograms on Large Data in MapReduce	109
<i>Jeffrey Jestes, Ke Yi, Feifei Li</i>	
Summarization and Matching of Density-Based Clusters in Streaming Environments	121
<i>Di Yang, Elke A. Rundensteiner, Matthew O. Ward</i>	
Multilingual Schema Matching for Wikipedia Infoboxes	133
<i>Thanh Nguyen, Viviane Moreira, Huong Nguyen, Hoa Nguyen, Juliana Freire</i>	

EXPERIMENTS AND ANALYSIS TRACK PAPERS

Controlling False Positives in Association Rule Mining	145
<i>Guimei Liu, Haojun Zhang, Limsoon Wong</i>	

NUMBER 3

RESEARCH TRACK PAPERS

PARIS: Probabilistic Alignment of Relations, Instances, and Schema	157
<i>Fabian M. Suchanek, Serge Abiteboul, Pierre Senellart</i>	
Answering Top-k Queries Over a Mixture of Attractive and Repulsive Dimensions	169
<i>Sayan Ranu, Ambuj K. Singh</i>	
PIQL: Success-Tolerant Query Processing in the Cloud	181
<i>Michael Armbrust, Kristal Curtis, Tim Kraska, Armando Fox, Michael J. Franklin, David A. Patterson</i>	
gSketch: On Query Estimation in Graph Streams	193
<i>Peixiang Zhao, Charu C. Aggarwal, Min Wang</i>	

Indexing the Earth Mover's Distance Using Normal Distributions	205
<i>Brian E. Rutenber, Ambuj K. Singh</i>	
Generating Exact- and Ranked Partially-Matched Answers to Questions in Advertisements	217
<i>Rani Qumsiyeh, Maria S. Pera, Yiu-Kai Ng</i>	
Size-1 Object Summaries for Relational Keyword Search	229
<i>Georgios J. Fakas, Zhi Cai, Nikos Mamoulis</i>	
REX: Explaining Relationships between Entity Pairs	241
<i>Lujun Fang, Anish Das Sarma, Cong Yu, Philip Bohannon</i>	
PASS-JOIN: A Partition-based Method for Similarity Joins	253
<i>Guoliang Li, Dong Deng, Jiannan Wang, Jianhua Feng</i>	
Relative Lempel-Ziv Factorization for Efficient Storage and Retrieval of Web Collections	265
<i>Christopher Hoobin, Simon J. Puglisi, Justin Zobel</i>	

NUMBER 4

RESEARCH TRACK PAPERS

Towards Cost-Effective Storage Provisioning for DBMSs	274
<i>Ning Zhang, Junichi Tatemura, Jignesh M. Patel, Hakan Hacigümüs</i>	
B+-tree Index Optimization by Exploiting Internal Parallelism of Flash-based Solid State Drives	286
<i>Hongchan Roh, Sanghyun Park, Sungho Kim, Mincheol Shin, Sang-Won Lee</i>	
High-Performance Concurrency Control Mechanisms for Main-Memory Databases	298
<i>Per-Åke Larson, Spyros Blanas, Cristian Diaconu, Craig Freedman, Jignesh M. Patel, Mike Zwilling</i>	
Capturing Topology in Graph Pattern Matching	310
<i>Shuai Ma, Yang Cao, Wenfei Fan, Jinpeng Huai, Tianyu Wo</i>	
Probabilistic Management of OCR Data using an RDBMS	322
<i>Arun Kumar, Christopher Ré</i>	
RTED: A Robust Algorithm for the Tree Edit Distance	334
<i>Mateusz Pawlik, Nikolaus Augsten</i>	
Putting Lipstick on Pig: Enabling Database-style Workflow Provenance	346
<i>Yael Amsterdamer, Susan B. Davidson, Daniel Deutch, Tova Milo, Julia Stoyanovich, Val Tannen</i>	
Relational Approach for Shortest Path Discovery over Large Graphs	358
<i>Jun Gao, Ruoming Jin, Jiashuai Zhou, Jeffrey Xu Yu, Xiao Jiang, Tengjiao Wang</i>	
Mining Flipping Correlations from Large Datasets with Taxonomies	370
<i>Marina Barsky, Sangkyum Kim, Tim Weninger, Jiawei Han</i>	
A Statistical Approach Towards Robust Progress Estimation	382
<i>Arnd Christian König, Bolin Ding, Surajit Chaudhuri, Vivek Narasayya</i>	

NUMBER 5

RESEARCH TRACK PAPERS

Relation Strength-Aware Clustering of Heterogeneous Information Networks with Incomplete Attributes	394
<i>Yizhou Sun, Charu C. Aggarwal, Jiawei Han</i>	

EXPERIMENTS AND ANALYSIS TRACK PAPERS

Shortest Path and Distance Queries on Road Networks: An Experimental Evaluation	406
<i>Lingkun Wu, Xiaokui Xiao, Dingxiang Deng, Gao Cong, Andy Diwen Zhu, Shuigeng Zhou</i>	

RESEARCH TRACK PAPERS

The Filter-Placement Problem and its Application to Minimizing Information Multiplicity	418
<i>Dóra Erdős, Vatche Ishakian, Andrei Lapets, Evimaria Terzi, Azer Bestavros</i>	
Bayesian Locality Sensitive Hashing for Fast Similarity Search	430
<i>Venu Satuluri, Srinivasan Parthasarathy</i>	

Fast and Exact Top-k Search for Random Walk with Restart	442
<i>Yasuhiro Fujiwara, Makoto Nakatsuji, Makoto Onizuka, Masaru Kitsuregawa</i>	
Densest Subgraph in Streaming and MapReduce	454
<i>Bahman Bahmani, Ravi Kumar, Sergei Vassilvitskii</i>	
Mining Attribute-Structure Correlated Patterns in Large Attributed Graphs	466
<i>Arlei Silva, Wagner Meira Jr., Mohammed J. Zaki</i>	
Semi-Automatic Index Tuning: Keeping DBAs in the Loop	478
<i>Karl Schnaitter, Neoklis Polyzotis</i>	
Aggregation in Probabilistic Databases via Knowledge Compilation	490
<i>Robert Fink, Larisa Han, Dan Olteanu</i>	

NUMBER 6

RESEARCH TRACK PAPERS

Stochastic Database Cracking: Towards Robust Adaptive Indexing in Main-Memory Column-Stores	502
<i>Felix Halim, Stratos Idreos, Panagiotis Karras, Roland H. C. Yap</i>	
An Adaptive Mechanism for Accurate Query Answering under Differential Privacy	514
<i>Chao Li, Gerome Miklau</i>	
SharedDB: Killing One Thousand Queries With One Stone	526
<i>Georgios Giannikis, Gustavo Alonso, Donald Kossmann</i>	
Pushing the Boundaries of Crowd-enabled Databases with Query-driven Schema Expansion	538
<i>Joachim Selke, Christoph Lofi, Wolf-Tilo Balke</i>	
A Bayesian Approach to Discovering Truth from Conflicting Sources for Data Integration	550
<i>Bo Zhao, Benjamin I. P. Rubinstein, Jim Gemmell, Jiawei Han</i>	
How to Price Shared Optimizations in the Cloud	562
<i>Prasang Upadhyaya, Magdalena Balazinska, Dan Suciu</i>	
Dense Subgraph Maintenance under Streaming Edge Weight Updates for Real-time Story Identification	574
<i>Albert Angel, Nick Koudas, Nikos Sarkas, Divesh Srivastava</i>	
ReStore: Reusing Results of MapReduce Jobs	586
<i>Iman Elghandour, Ashraf Aboulnaga</i>	

NUMBER 7

RESEARCH TRACK PAPERS

PerfXplain: Debugging MapReduce Job Performance	598
<i>Nodira Khousainova, Magdalena Balazinska, Dan Suciu</i>	
Uncertain Centroid based Partitional Clustering of Uncertain Data	610
<i>Francesco Gullo, Andrea Tagarelli</i>	
Scalable K-Means++	622
<i>Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar, Sergei Vassilvitskii</i>	
Querying Schemas With Access Restrictions	634
<i>Michael Benedikt, Pierre Bourhis, Clemens Ley</i>	
Definition, Detection, and Recovery of Single-Page Failures, a Fourth Class of Database Failures	646
<i>Goetz Graefe, Harumi Kuno</i>	

Volume 2

Concurrency Control for Adaptive Indexing	656
<i>Goetz Graefe, Felix Halim, Stratos Idreos, Harumi Kuno, Stefan Manegold</i>	

EXPERIMENTS AND ANALYSIS TRACK PAPERS

Comments on "Stack-based Algorithms for Pattern Matching on DAGs"	668
<i>Qiang Zeng, Hai Zhuge</i>	

An Analysis of Structured Data on the Web	680
<i>Nilesh Dalvi, Ashwin Machanavajjhala, Bo Pang</i>	

NUMBER 8

RESEARCH TRACK PAPERS

Shortest Path Computation with No Information Leakage	692
<i>Kyriakos Mouratidis, Man Lung Yiu</i>	
V-SMART-Join: A Scalable MapReduce Framework for All-Pair Similarity Joins of Multisets and Vectors	704
<i>Ahmed Metwally, Christos Faloutsos</i>	
Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud	716
<i>Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, Joseph M. Hellerstein</i>	
Adding Logical Operators to Tree Pattern Queries on Graph-Structured Data	728
<i>Qiang Zeng, Xiaorui Jiang, Hai Zhuge</i>	
Learning Semantic String Transformations from Examples	740
<i>Rishabh Singh, Sumit Gulwani</i>	
Cologne: A Declarative Distributed Constraint Optimization Platform	752
<i>Changbin Liu, Lu Ren, Boon Thau Loo, Yun Mao, Prithwish Basu</i>	
Optimizing I/O for Big Array Analytics	764
<i>Yi Zhang, Jun Yang</i>	
Probabilistically Bounded Staleness for Practical Partial Quorums	776
<i>Peter Bailis, Shivaram Venkataraman, Michael J. Franklin, Joseph M. Hellerstein, Ion Stoica</i>	

NUMBER 9

RESEARCH TRACK PAPERS

Efficient Subgraph Matching on Billion Node Graphs	788
<i>Zhao Sun, Hongzhi Wang, Haixun Wang, Bin Shao, Jianzhong Li</i>	
Efficient Subgraph Similarity Search on Large Probabilistic Graph Databases	800
<i>Ye Yuan, Guoren Wang, Lei Chen, Haixun Wang</i>	
Truss Decomposition in Massive Networks	812
<i>Jia Wang, James Cheng</i>	
SEAL: Spatio-Textual Similarity Search	824
<i>Ju Fan, Guoliang Li, Lizhu Zhou, Shanshan Chen, Jun Hu</i>	
On The Spatiotemporal Burstiness of Terms	836
<i>Theodoros Lappas, Marcos R. Vieira, Dimitrios Gunopulos, Vassilis J. Tsotras</i>	
Efficient Reachability Query Evaluation in Large Spatiotemporal Contact Datasets	848
<i>Houtan Shirani-Mehr, Farnoush Banaei Kashani, Cyrus Shahabi</i>	
Boosting Moving Object Indexing through Velocity Partitioning	860
<i>Thi Nguyen, Zhen He, Rui Zhang, Phillip Ward</i>	
Type-Based Detection of XML Query-Update Independence	872
<i>Nicole Bidoit-Tollu, Dario Colazzo, Federico Ulliana</i>	
Minuet: A Scalable Distributed Multiversion B-Tree	884
<i>Benjamin Sowell, Wojciech Golab, Mehul A. Shah</i>	
Challenging the Long Tail Recommendation	896
<i>Hongzhi Yin, Bin Cui, Jing Li, Junjie Yao, Chen Chen</i>	

NUMBER 10

RESEARCH TRACK PAPERS

Answering Table Queries on the Web using Column Keywords	908
<i>Rakesh Pimplikar, Sunita Sarawagi</i>	
Efficient Verification of Web-Content Searching Through Authenticated Web Crawlers	920
<i>Michael T. Goodrich, Duy Nguyen, Olga Ohrimenko, Charalampos Papamanthou, Roberto Tamassia, Nikos Triandopoulos, Cristina Videira Lopes</i>	

SODA: Generating SQL for Business Users	932
<i>Lukas Blunschi, Claudio Jossen, Donald Kossmann, Magdalini Mori, Kurt Stockinger</i>	
Privacy Preservation by Disassociation	944
<i>Manolis Terrovitis, John Liagouris, Nikos Mamoulis, Spiros Skiadopoulos</i>	
Supercharging Recommender Systems using Taxonomies for Learning User Purchase Behavior	956
<i>Bhargav Kanagal, Amr Ahmed, Sandeep Pandey, Vanja Josifovski, Jeff Yuan, Lluís Garcia-Pueyo</i>	
DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views	968
<i>Yanif Ahmad, Oliver Kennedy, Christoph Koch, Milos Nikolic</i>	
Real Time Discovery of Dense Clusters in Highly Dynamic Graphs: Identifying Real World Events in Highly Dynamic Environments	980
<i>Manoj K. Agarwal, Krithi Ramamritham, Manish Bhide</i>	
Sketch-based Querying of Distributed Sliding-Window Data Streams	992
<i>Odysseas Papapetrou, Minos Garofalakis, Antonios Deligiannakis</i>	
LogBase: A Scalable Log-structured Database System in the Cloud	1004
<i>Hoang Tam Vo, Sheng Wang, Divyakant Agrawal, Gang Chen, Beng Chin Ooi</i>	
Efficient Processing of k Nearest Neighbor Joins using MapReduce	1016
<i>Wei Lu, Yanyan Shen, Su Chen, Beng Chin Ooi</i>	
Early Accurate Results for Advanced Analytics on MapReduce	1028
<i>Nikolay Laptev, Kai Zeng, Carlo Zaniolo</i>	
CDAS: A Crowdsourcing Data Analytics System	1040
<i>Xuan Liu, Meiyu Lu, Beng Chin Ooi, Yanyan Shen, Sai Wu, Meihui Zhang</i>	
Mining Statistically Significant Substrings using the Chi-Square Statistic	1052
<i>Mayank Sachan, Arnab Bhattacharya</i>	
Massively Parallel Sort-Merge Joins in Main Memory Multi-Core Database Systems	1064
<i>Martina-Cezara Albutiu, Alfons Kemper, Thomas Neumann</i>	
hStorage-DB: Heterogeneity-aware Data Management to Exploit the Full Capability of Hybrid Storage Systems	1076
<i>Tian Luo, Rubao Lee, Michael Mesnier, Feng Chen, Xiaodong Zhang</i>	

NUMBER 11

RESEARCH TRACK PAPERS

A Scalable Algorithm for Maximizing Range Sum in Spatial Databases	1088
<i>Dong-Wan Choi, Chin-Wan Chung, Yufei Tao</i>	
Spatial Queries with Two kNN Predicates	1100
<i>Ahmed M. Aly, Walid G. Aref, Mourad Ouzzani</i>	
Optimal Algorithms for Crawling a Hidden Database in the Web	1112
<i>Cheng Sheng, Nan Zhang, Yufei Tao, Xin Jin</i>	
Diversifying Top-K Results	1124
<i>Lu Qin, Jeffrey Xu Yu, Lijun Chang</i>	
Keyword-aware Optimal Route Search	1136
<i>Xin Cao, Lisi Chen, Gao Cong, Xiaokui Xiao</i>	
Answering Queries using Views over Probabilistic XML: Complexity and Tractability	1148
<i>Bogdan Cautis, Evgeny Kharlamov</i>	
Probabilistic Databases with MarkoViews	1160
<i>Abhay Jha, Dan Suciu</i>	
The Complexity of Social Coordination	1172
<i>Konstantinos Mamouras, Sigal Oren, Lior Seeman, Lucja Kot, Johannes Gehrke</i>	
Efficient Multi-way Theta-Join Processing Using MapReduce	1184
<i>Xiaofei Zhang, Lei Chen, Min Wang</i>	
Stubby: A Transformation-based Optimizer for MapReduce Workflows	1196
<i>Harold Lim, Herodotos Herodotou, Shiyath Babu</i>	
Labeling Workflow Views with Fine-Grained Dependencies	1208
<i>Zhuowei Bao, Susan B. Davidson, Tova Milo</i>	
Fundamentals of Order Dependencies	1220
<i>Jaroslav Szlichta, Parke Godfrey, Jarek Gryz</i>	
FDB: A Query Engine for Factorised Relational Databases	1232
<i>Nurzhan Bakibayev, Dan Olteanu, Jakub Závodný</i>	
Optimization of Analytic Window Functions	1244
<i>Yu Cao, Chee-Yong Chan, Jie Li, Kian-Lee Tan</i>	

Opening the Black Boxes in Data Flow Optimization	1256
<i>Fabian Hueske, Mathias Peters, Matthias Sax, Astrid Rheinländer, Rico Bergmann, Aljoscha Krettek, Kostas Tzoumas</i>	
Spinning Fast Iterative Data Flows	1268
<i>Stephan Ewen, Kostas Tzoumas, Moritz Kaufmann, Volker Markl</i>	
REX: Recursive, Delta-Based Data-Centric Computation	1280
<i>Svilen R. Mihaylov, Zachary G. Ives, Sudipto Guha</i>	
K-Reach: Who is in Your Small World	1292
<i>James Cheng, Zechao Shang, Hong Cheng, Haixun Wang, Jeffrey Xu Yu</i>	
Performance Guarantees for Distributed Reachability Queries	1304
<i>Wenfei Fan, Xin Wang, Yinghui Wu</i>	
Efficient Indexing and Querying over Syntactically Annotated Trees	1316
<i>Pirooz Chubak, Davood Rafiei</i>	
Queries with Guarded Negation	1328
<i>Vince Barany, Balder Ten Cate, Martin Otto</i>	

Volume 3

PrivBasis: Frequent Itemset Mining with Differential Privacy	1340
<i>Ninghui Li, Wahbeh Qardaji, Dong Su, Jianneng Cao</i>	
Low-Rank Mechanism: Optimizing Batch Queries under Differential Privacy	1352
<i>Ganzhao Yuan, Zhenjie Zhang, Marianne Winslett, Xiaokui Xiao, Yin Yang, Zhifeng Hao</i>	
Functional Mechanism: Regression Analysis under Differential Privacy	1364
<i>Jun Zhang, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, Marianne Winslett</i>	
Injecting Uncertainty in Graphs for Identity Obfuscation	1376
<i>Paolo Boldi, Francesco Bonchi, Aris Gionis, Tamir Tassa</i>	
Publishing Microdata with a Robust Privacy Guarantee	1388
<i>Jianneng Cao, Panagiotis Karras</i>	
Measuring Two-Event Structural Correlations on Graphs	1400
<i>Ziyu Guan, Xifeng Yan, Lance M. Kaplan</i>	
Ranking Large Temporal Data	1412
<i>Jeffrey Jestes, Jeff M. Phillips, Feifei Li, Mingwang Tang</i>	
Compacting Transactional Data in Hybrid OLTP & OLAP Databases	1424
<i>Florian Funke, Alfons Kemper, Thomas Neumann</i>	
Processing a Trillion Cells per Mouse Click	1436
<i>Alexander Hall, Olaf Bachmann, Robert Büssow, Silviu Ganceanu, Marc Nunkesser</i>	
OLTP on Hardware Islands	1447
<i>Danica Porobic, Ippokratis Pandis, Miguel Branco, Pinar Tözün, Anastasia Ailamaki</i>	
Serializability, not Serial: Concurrence Control and Availability in Multi-Datacenter Datastores	1459
<i>Stacy Patterson, Aaron J. Elmore, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi</i>	
Automatic Partitioning of Database Applications	1471
<i>Alvin Cheung, Owen Arden, Samuel Madden, Andrew C. Myers</i>	
CrowdER: Crowdsourcing Entity Resolution	1483
<i>Jiannan Wang, Tim Kraska, Michael J. Franklin, Jianhua Feng</i>	
Whom to Ask? Jury Selection for Decision Making Tasks on Micro-blog Services	1495
<i>Caleb Chen Cao, Jieying She, Yongxin Tong, Lei Chen</i>	
ALAE: Accelerating Local Alignment with Affine Gap Exactly in Biosequence Databases	1507
<i>Xiaochun Yang, Honglei Liu, Bin Wang</i>	
sDTW: Computing DTW Distances using Locally Relevant Constraints based on Salient Feature Alignments	1519
<i>K. Selçuk Candan, Rosaria Rossini, Maria Luisa Sapino, Xiaolan Wang</i>	
SCOUT: Prefetching for Latent Feature Following Queries	1531
<i>Farhan Tauheed, Thomas Heinis, Felix Shürmann, Henry Markram, Anastasia Ailamaki</i>	
Accelerating Pathology Image Data Cross-Comparison on CPU-GPU Hybrid Systems	1543
<i>Kaibo Wang, Yin Huai, Rubao Lee, Fusheng Wang, Xiaodong Zhang, Joel H. Saltz</i>	
Robust Estimation of Resource Consumption for SQL Queries using Statistical Techniques	1555
<i>Jiexing Li, Arnd Christian König, Vivek Narasayya, Surajit Chaudhuri</i>	
Who Tags What? An Analysis Framework	1567
<i>Mahashweta Das, Saravanan Thirumuruganathan, Sihem Amer-Yahia, Gautam Das, Cong Yu</i>	
A Generic Framework for Efficient and Effective Subsequence Retrieval	1579
<i>Haohan Zhu, George Kollios, Vassilis Athitsos</i>	

Only Aggressive Elephants are Fast Elephants	1591
<i>Jens Dittrich, Jorge-Arnulfo Quiané-Ruiz, Stefan Richter, Stefan Schuh, Alekh Jindal, Jörg Schäd</i>	
Multiple Location Profiling for Users and Relationships from Social Network and Content	1603
<i>Rui Li, Shengjie Wang, Kevin Chen-Chuan Chang</i>	
Flash-based Extended Cache for Higher Throughput and Faster Recovery	1615
<i>Woon-Hak Kang, Sang-Won Lee, Bongki Moon</i>	
Don't Thrash: How to Cache Your Hash on Flash	1627
<i>Michael A. Bender, Martin Farach-Colton, Rob Johnson, Russell Kraner, Bradley C. Kuszmaul, Dzejla Medjedovic, Pablo Montes, Pradeep Shetty, Richard P. Spillane, Erez Zadok</i>	
Learning Expressive Linkage Rules using Genetic Programming	1638
<i>Robert Isele, Christian Bizer</i>	

EXPERIMENTS AND ANALYSIS TRACK PAPERS

Mining Frequent Itemsets over Uncertain Databases	1650
<i>Yongxin Tong, Lei Chen, Yurong Cheng, Philip S. Yu</i>	
Uncertain Time-Series Similarity: Return to the Basics	1662
<i>Michele Dallachiesa, Besmira Nushi, Katsiaryna Mirylenka, Themis Palpanas</i>	
Statistical Distortion: Consequences of Data Cleaning	1674
<i>Tamraparni Dasu, Ji Meng Loh</i>	
Towards Energy-Efficient Database Cluster Design	1684
<i>Willis Lang, Stavros Harizopoulos, Jignesh M. Patel, Mehul A. Shah, Dimitris Tsirogiannis</i>	

NUMBER 12

KEYNOTES

Data Management on the Spatial Web	1696
<i>Christian S. Jensen</i>	
Data Analytics Opportunities in a Smarter Planet	1697
<i>Brenda Dietrich</i>	
Challenges in Economic Massive Content Storage and Management (MCSAM) in the Era of Self-Organizing, Self-Expanding and Self-Linking Data Clusters	1698
<i>Kenan Sahin</i>	

10-YEAR BEST PAPER AWARD

Approximate Frequency Counts over Data Streams	1699
<i>Gurmeet Singh Manku, Rajeev Motwani</i>	

INDUSTRIAL, APPLICATIONS, AND EXPERIENCE TRACK PAPERS

The MADlib Analytics Library or MAD Skills, the SQL	1700
<i>Joe Hellerstein, Christopher Ré, Florian Schoppmann, Daisy Zhe Wang, Eugene Fratkin, Aleksander Gorajek, Kee Siong Ng, Caleb Welton, Xixuan Feng, Kun Li, Arun Kumar</i>	
Can the Elephants Handle the NoSQL Onslaught?	1712
<i>Avrilia Floratou, Nikhil Teletia, David J. Dewitt, Jignesh M. Patel, Donghui Zhang</i>	
Solving Big Data Challenges for Enterprise Application Performance Management	1724
<i>Tilmann Rabl, Mohammad Sadoghi, Hans-Arno Jacobsen, Sergio Gómez-Villamor, Victor Muntés-Mulero, Serge Mankowski</i>	
M3R: Increased Performance for In-memory Hadoop Jobs	1736
<i>Avraham Shinnar, David Cunningham, Benjamin Herta, Vijay Saraswat</i>	
A Storage Advisor for Hybrid-Store Databases	1748
<i>Philipp Rösch, Lars Dannecker, Gregor Hackenbroich, Franz Faerber</i>	
From Cooperative Scans to Predictive Buffer Management	1759
<i>Michał Swiatkowski, Peter Boncz, Marcin Zukowski</i>	
The Unified Logging Infrastructure for Data Analytics at Twitter	1771
<i>George Lee, Jimmy Lin, Chuang Liu, Andrew Lorek, Dmitriy Ryaboy</i>	

Transaction Log Based Application Error Recovery and Point In-Time Query	1781
<i>Tomas Talius, Robin Dhamankar, Andrei Dumitrache, Hanuma Kodavalla</i>	
The Vertica Analytic Database: C-Store 7 Years Later	1790
<i>Andrew Lamb, Matt Fuller, Ramakrishna Varadarajan, Nga Tran, Ben Vandier, Lyric Doshi, Chuck Bear</i>	
Interactive Analytical Processing in Big Data Systems: A Cross-Industry Study of MapReduce Workloads	1802
<i>Yanpei Chen, Sara Alspaugh, Randy Katz</i>	
Muppet: MapReduce-Style Processing of Fast Data	1814
<i>Wang Lam, Lu Liu, S. T. S. Prasad, Anand Rajaraman, Zoheb Vacheri, Anhai Doan</i>	
Building User-defined Runtime Adaptation Routines for Stream Processing Applications	1826
<i>Gabriela Jacques-Silva, Bugra Gedik, Rohit Wagle, Kun-Lung Wu, Vibhore Kumar</i>	
MOIST: A Scalable and Parallel Moving Object Indexer with School Tracking	1838
<i>Junchen Jiang, Hongji Bao, Edward Y. Chang, Yuqian Li</i>	
Serializable Snapshot Isolation in PostgreSQL	1850
<i>Dan R. K. Ports, Kevin Grittner</i>	
Exploiting Evidence from Unstructured Data to Enhance Master Data Management	1862
<i>Karin Murthy, Prasad M. Deshpande, Atreyee Dey, Ramanujam Halasipuram, Mukesh Mohania, P. Deepak, Jennifer Reed, Scott Schumacher</i>	
Avatara: OLAP for Web-scale Analytics Products	1874
<i>Lili Wu, Roshan Sumbaly, Chris Riccomini, Gordon Koo, Hyung Jin Kim, Jay Kreps, Sam Shah</i>	

DEMONSTRATION TRACK PAPERS

Dedoop: Efficient Deduplication with Hadoop	1878
<i>Lars Kolb, Andreas Thor, Erhard Rahm</i>	
MapReduce-based Dimensional ETL Made Easy	1882
<i>Xiufeng Liu, Christian Thomsen, Torben Bach Pedersen</i>	
CloudVista: Interactive and Economical Visual Cluster Analysis for Big Data in the Cloud	1886
<i>Huiqi Xu, Zhen Li, Shumin Guo, Keke Chen</i>	
Myriad: Scalable and Expressive Data Generation	1890
<i>Alexander Alexandrov, Kostas Tzoumas, Volker Markl</i>	
A Demonstration of DBWipes: Clean as You Query	1894
<i>Eugene Wu, Samuel Madden, Michael Stonebraker</i>	
ASTERIX: An Open Source System for "Big Data" Management and Analysis	1898
<i>Sattam Alsubaiee, Yasser Altowim, Hotham Altwaijry, Alexander Behm, Vinayak Borkar, Yingyi Bu, Michael Carey, Raman Grover, Zachary Heilbron, Young-Seok Kim, Chen Li, Nicola Onose, Pouria Pirzadeh, Rares Vernica, Jian Wen</i>	
Blink and It's Done: Interactive Queries on Very Large Data	1902
<i>Sameer Agarwal, Aurojit Panda, Barzan Mozafari, Anand P. Iyer, Samuel Madden, Ion Stoica</i>	
Massive Genomic Data Processing and Deep Analysis	1906
<i>Abhishek Roy, Yanlei Diao, Evan Mauceli, Yiping Shen, Bai-Lin Wu</i>	
MonetDB/DataCell: Online Analytics in a Streaming Column-Store	1910
<i>Erietta Liarou, Stratos Idreos, Stefan Manegold, Martin Kersten</i>	
SWORS: A System for the Efficient Retrieval of Relevant Spatial Web Objects	1914
<i>Xin Cao, Gao Cong, Christian S. Jensen, Jun Jie Ng, Beng Chin Ooi, Nhan-Tue Phan, Dingming Wu</i>	
CyLog/Crowd4U: A Declarative Platform for Complex Data-centric Crowdsourcing	1918
<i>Atsuyuki Morishima, Norihide Shinagawa, Tomomi Mitsuishi, Hideto Aoki, Shun Fukusumi</i>	
Exploiting Database Similarity Joins for Metric Spaces	1922
<i>Yasin N. Silva, Spencer Pearson</i>	
Stethoscope: A Platform for Interactive Visual Analysis of Query Execution Plans	1926
<i>Mrunal Gawade, Martin Kersten</i>	
Hum-a-song: A Subsequence Matching with Gaps-Range-Tolerances Query-By-Humming System	1930
<i>Alexios Kotsifakos, Panagiotis Papapetrou, Jaakko Hollmén, Dimitrios Gunopoulos, Vassilis Athitsos, George Kollios</i>	
SkewTune in Action: Mitigating Skew in MapReduce Applications	1934
<i>Yongchul Kwon, Magdalena Balazinska, Bill Howe, Jerome Rolia</i>	
Playful Query Specification with DataPlay	1938
<i>Azza Abouzied, Joseph M. Hellerstein, Avi Silberschatz</i>	
NoDB in Action: Adaptive Query Processing on Raw Data	1942
<i>Ioannis Alagiannis, Renata Borovica, Miguel Branco, Stratos Idreos, Anastasia Ailamaki</i>	
Complex Preference Queries Supporting Spatial Applications for User Groups	1946
<i>Florian Wenzel, Markus Endres, Stefan Mandl, Werner Kießling</i>	

Demonstration of the FDB Query Engine for Factorised Databases	1950
<i>Nurzhan Bakibayev, Dan Olteanu, Jakub Závodný</i>	
PET: Reducing Database Energy Cost via Query Optimization	1954
<i>Zichen Xu, Yi-Cheng Tu, Xiaorui Wang</i>	
SPAM: A SPARQL Analysis and Manipulation Tool	1958
<i>Andrés Letelier, Jorge Pérez, Reinhard Pichler, Sebastian Skritek</i>	
QueryMarket Demonstration: Pricing for Online Data Markets	1962
<i>Paraschos Koutris, Prasang Upadhyaya, Magdalena Balazinska, Bill Howe, Dan Suciu</i>	
DISKS: A System for Distributed Spatial Group Keyword Search on Road Networks	1966
<i>Iqiang Luo, Yifeng Luo, Shuigeng Zhou, Gao Cong, Jihong Guan</i>	
WETSUIT: An Efficient Mashup Tool for Searching and Fusing Web Entities	1970
<i>Stefan Endrullis, Andreas Thor, Erhard Rahm</i>	
Model-based Integration of Past & Future in TimeTravel	1974
<i>Mohamed E. Khalefa, Ulrike Fischer, Torben Bach Pedersen, Wolfgang Lehner</i>	
DrillBeyond: Enabling Business Analysts to Explore the Web of Open Data	1978
<i>Julian Eberius, Maik Thiele, Katrin Braunschweig, Wolfgang Lehner</i>	
Discovering and Exploring Relations on the Web	1982
<i>Ndapandula Nakashole, Gerhard Weikum, Fabian Suchanek</i>	
MapRat: Meaningful Explanation, Interactive Exploration and Geo-Visualization of Collaborative Ratings	1986
<i>Saravanan Thirumuruganathan, Mahashweta Das, Shrikant Desai, Sihem Amer-Yahia, Gautam Das, Cong Yu</i>	
Deco: A System for Declarative Crowdsourcing	1990
<i>Hyunjung Park, Richard Pang, Aditya Parameswaran, Hector Garcia-Molina, Neoklis Polyzotis, Jennifer Widom</i>	
Developing and Analyzing XSDs through BonXai	1994
<i>Wim Martens, Frank Neven, Matthias Niewerth, Thomas Schwentick</i>	
InfoPuzzle: Exploring Group Decision Making in Mobile Peer-to-Peer Databases	1998
<i>Aaron J. Elmore, Sudipto Das, Divyakant Agrawal, Amr El Abbadi</i>	
Manage and Query Generic Moving Objects in SECONDO	2002
<i>Jianqiu Xu, Ralf Hartmut Güting</i>	
Chronos: Facilitating History Discovery by Linking Temporal Records	2006
<i>Pei Li, Haidong Wang, Christina Tziviskou, Xin Luna Dong, Xiaoguang Liu, Andrea Maurino, Divesh Srivastava</i>	
TELEIOS: A Database-Powered Virtual Earth Observatory	2010
<i>Manolis Koubarakis, Kostis Kyzirakos, Manos Karpathiotakis, Charalampos Nikolaou, Stavros Vassos, George Garbis, Michael Sioutis, Konstantina Bereta, Dimitrios Michail, Charalampos Kontoes, Ioannis Papoutsis, Themos Herekakis, Stefan Manegold, Martin Kersten, Milena Ivanova, Holger Pirk, Ying Zhang, Mihai Datcu, Gottfried Schwarz, Corneliu Dumitru, Daniela Espinoza Molina, Katrin Molch, Ugo Di Giammatteo, Manuela Sagona, Sergio Perelli, Thorsten Reitz, Eva Klien, Robert Gregor</i>	

TUTORIALS

Efficient Big Data Processing in Hadoop MapReduce	2014
<i>Jens Dittrich, Jorge-Arnulfo Quiané-Ruiz</i>	
MapReduce Algorithms for Big Data Analysis	2016
<i>Kyuseok Shim</i>	
Entity Resolution: Theory, Practice & Open Challenges	2018
<i>Lise Getoor, Ashwin Machanavajjhala</i>	
I/O Characteristics of NoSQL Databases	2020
<i>Jiri Schindler</i>	
Mining Knowledge from Interconnected Data: A Heterogeneous Information Network Analysis Approach	2022
<i>Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu</i>	
Understanding and Managing Cascades on Large Graphs	2024
<i>B. Aditya Prakash, Christos Faloutsos</i>	
Interoperability in eHealth Systems (Invited Tutorial)	2026
<i>Asuman Dogac</i>	
Secure and Privacy-Preserving Data Services in the Cloud: A Data Centric View	2028
<i>Divyakant Agrawal, Amr El Abbadi, Shiyuan Wang</i>	
Graph Synopses, Sketches, and Streams: A Survey	2030
<i>Sudipto Guha, Andrew McGregor</i>	

PANELS

Challenges and Opportunities with Big Data 2032
Alexandros Labrinidis, H. V. Jagadish

Social Networks and Mobility in the Cloud 2034
Amr El Abbadi, Mohamed F. Mokbel

Author Index