

2008 IEEE 24th International Conference on Data Engineering

**Cancun, Mexico
7-12 April 2008**

Pages 1-545



IEEE Catalog Number:
ISBN 10:
ISBN 13:

CFP08026-PRT
1-4244-1836-4
978-1-4244-1836-7

Table of Contents

Human Computation	1
<i>Luis von Ahn</i>	
The Database Architecture Jigsaw Puzzle	3
<i>Martin L. Kersten</i>	
PhotoSpread: A Spreadsheet for Managing Photos	5
<i>Hector Garcia-Molina</i>	
Declarative, Domain-Specific Languages --- Elegant Simplicity or a Hammer in Search of a Nail?	7
<i>Samuel Madden, Johannes Gehrke</i>	
Cloud Computing --- Was Thomas Watson Right After All?	8
<i>Raghu Ramakrishnan</i>	
Scientific Data Management: An Orphan in the Database Community?	9
<i>Randal Burns</i>	
Muse: Mapping Understanding and deSign by Example	10
<i>Bogdan Alexe, Laura Chiticariu, Renee J. Miller, Wang-Chiew Tan</i>	
Usage-Based Schema Matching	20
<i>Hazem Elmeleegy, Mourad Ouzzani, Ahmed K. Elmagarmid</i>	
Clip: A Visual Language for Explicit Schema Mappings	30
<i>Alessandro Raffio, Daniele Braga, Stefano Ceri, Paolo Papotti, Mauricio A. Hernandez</i>	
Transformation-Based Framework for Record Matching	40
<i>Arvind Arasu, Surajit Chaudhuri, Raghav Kaushik</i>	
Querying Data Under Access Limitations	50
<i>Andrea Cali, Davide Martinenghi</i>	
Constant-Time Query Processing	60
<i>Vijayshankar Raman, Garret Swart, Lin Qiao, Frederick Reiss, Vijay Dialani, Donald Kossmann, Inderpal Narang, Richard Sidle</i>	
A Hybrid Prediction Model for Moving Objects	70
<i>Hoyoung Jeung, Qing Liu, Heng Tao Shen, Xiaofang Zhou</i>	
A Generic Framework for Continuous Motion Pattern Query Evaluation	80
<i>Petko Bakalov, Vassilis J. Tsotras</i>	
Querying Complex Spatio-Temporal Sequences in Human Motion Databases	90
<i>Yueguo Chen, Shouxu Jiang, Beng Chin Ooi, Anthony K.H. Tung</i>	
Common Influence Join: A Natural Join Operation for Spatial Pointsets	100
<i>Man Lung Yiu, Nikos Mamoulis, Panagiotis Karras</i>	
Matching Schemas in Online Communities: A Web 2.0 Approach	110
<i>Robert McCann, Warren Shen, AnHai Doan</i>	
Validating Multi-Column Schema Matchings by Type	120
<i>Bing Tian Dai, Nick Koudas, Divesh Srivastava, Anthony K.H. Tung, Suresh Venkatasubramanian</i>	
CARE: Finding Local Linear Correlations in High Dimensional Data	130
<i>Xiang Zhang, Feng Pan, Wei Wang</i>	
Trajectory Outlier Detection: A Partition-and-Detect Framework	140
<i>Jae-Gil Lee, Jiawei Han, Xiaolei Li</i>	
A Framework for Clustering Uncertain Data Streams	150
<i>Charu C. Aggarwal, Philip S. Yu</i>	

Table of Contents

Mining Approximate Order Preserving Clusters in the Presence of Noise	160
<i>Mengsheng Zhang, Wei Wang, Jinze Liu</i>	
Direct Discriminative Pattern Mining for Effective Classification	169
<i>Hong Cheng, Xifeng Yan, Jiawei Han, Philip S. Yu</i>	
Verifying and Mining Frequent Patterns from Large Windows Over Data Streams	179
<i>Barzan Mozafari, Hetal Thakkar, Carlo Zaniolo</i>	
On Signatures for Communication Graphs	189
<i>Graham Cormode, Flip Korn, S. Muthukrishnan, Yihua Wu</i>	
Robust Stratified Sampling Plans for Low Selectivity Queries	199
<i>Shantanu Joshi, Christopher Jermaine</i>	
Training Linear Discriminant Analysis in Linear Time	209
<i>Deng Cai, Xiaofei He, Jiawei Han</i>	
Parallel Evaluation of Composite Aggregate Queries	218
<i>Lei Chen, Christopher Olston, Raghu Ramakrishnan</i>	
Efficient Computation of Diverse Query Results	228
<i>Erik Vee, Utkarsh Srivastava, Jayavel Shanmugasundaram, Prashant Bhat, Sihem Amer Yahia</i>	
Toward Expressive and Scalable Sponsored Search Auctions	237
<i>David J. Martin, Johannes Gehrke, Joseph Y. Halpern</i>	
Lattice Histograms: A Resilient Synopsis Structure	247
<i>Panagiotis Karras, Nikos Mamoulis</i>	
Efficient Merging and Filtering Algorithms for Approximate String Searches	257
<i>Chen Li, Jiaheng Lu, Yiming Lu</i>	
Fast Indexes and Algorithms for Set Similarity Selection Queries	267
<i>Marios Hadjieleftheriou, Amit Chandel, Nick Koudas, Divesh Srivastava</i>	
Privacy: Theory Meets Practice on the Map	277
<i>Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, Lars Vilhuber</i>	
Auditing SQL Queries	287
<i>Rajeev Motwani, Shubha U. Nabar, Dilys Thomas</i>	
Optimal Boolean Matrix Decomposition: Application to Role Engineering	297
<i>Haibing Lu, Jaideep Vaidya, Vijayalakshmi Atluri</i>	
Efficient Similarity Search Using the Earth Mover's Distance for Large Multimedia Databases	307
<i>Ira Assent, Marc Wichterich, Tobias Meisen, Thomas Seidl</i>	
Similarity Search in Arbitrary Subspaces Under L_p-Norm	317
<i>Xiang Lian, Lei Chen</i>	
Nearest Neighbor Retrieval Using Distance-Based Hashing	327
<i>Vassilis Athitsos, Michalis Potamias, Panagiotis Papapetrou, George Kollios</i>	
On Monitoring the Top-k Unsafe Places	337
<i>Donghui Zhang, Yang Du, Ling Hu</i>	
Compact Similarity Joins	346
<i>Brent Bryan, Frederick Eberhardt, Christos Faloutsos</i>	
Standing Out in a Crowd: Selecting Attributes for Maximum Visibility	356
<i>Muhammed Miah, Gautam Das, Vagelis Hristidis, Heikki Mannila</i>	
SpaceTwist: Managing the Trade-Offs Among Location Privacy, Query Performance, and Query Accuracy in Mobile Services	366
<i>Man Lung Yiu, Christian S. Jensen, Xuegang Huang, Hua Lu</i>	

Table of Contents

Never Walk Alone: Uncertainty for Anonymity in Moving Objects Databases	376
<i>Osman Abul, Francesco Bonchi, Mirco Nanni</i>	
On Unifying Privacy and Uncertain Data Models	386
<i>Charu C. Aggarwal</i>	
Efficient Construction of Compact Shedding Filters for Data Stream Processing	396
<i>Bugra Gedik, Kun-Lung Wu, Philip S. Yu</i>	
A Security Punctuation Framework for Enforcing Access Control on Streaming Data	406
<i>Rimma V. Nehme, Elke A. Rundensteiner, Elisa Bertino</i>	
Randomized Synopses for Query Assurance on Data Streams	416
<i>Ke Yi, Feifei Li, Marios Hadjieleftheriou, George Kollios, Divesh Srivastava</i>	
LOCUST: An Online Analytical Processing Framework for High Dimensional Classification of Data Streams	426
<i>Charu C. Aggarwal, Philip S. Yu</i>	
Dynamic Materialization of Query Views for Data Warehouse Workloads	436
<i>Thomas Phan, Wen-Syan Li</i>	
Injector: Mining Background Knowledge for Data Anonymization	446
<i>Tiancheng Li, Ninghui Li</i>	
RiTE: Providing On-Demand Data for Right-Time Data Warehousing	456
<i>Christian Thomsen, Torben Bach Pedersen, Wolfgang Lehner</i>	
Automatic Extraction of Useful Facet Hierarchies from Text Databases	466
<i>Wisam Dakka, Panagiotis G. Ipeirotis</i>	
A Sampling-Based Approach to Information Recovery	476
<i>Junyi Xie, Jun Yang, Yuguo Chen, Haixun Wang, Philip S. Yu</i>	
Mining (Social) Network Graphs to Detect Random Link Attacks	486
<i>Nisheeth Shrivastava, Anirban Majumder, Rajeev Rastogi</i>	
A Hybrid Approach to Private Record Linkage	496
<i>Ali Inan, Murat Kantarcioglu, Elisa Bertino, Monica Scannapieco</i>	
Preserving Privacy in Social Networks Against Neighborhood Attacks	506
<i>Bin Zhou, Jian Pei</i>	
Increasing the Expressivity of Conditional Functional Dependencies Without Extra Complexity	516
<i>Loreto Bravo, Wenfei Fan, Floris Geerts, Shuai Ma</i>	
Efficient Constraint Monitoring Using Adaptive Thresholds	526
<i>Srinivas Kashyap, Jeyashankher Ramamirtham, Rajeev Rastogi, Pushpraj Shukla</i>	
Dominant Graph: An Efficient Indexing Structure to Answer Top-k Queries	536
<i>Lei Zou, Lei Chen</i>	
Parallel Distributed Processing of Constrained Skyline Queries by Filtering	546
<i>Bin Cui, Hua Lu, Quanqing Xu, Lijiang Chen, Yafei Dai, Yongluan Zhou</i>	
Skyline Query Processing for Incomplete Data	556
<i>Mohamed E. Khalefa, Mohamed F. Mokbel, Justin J. Levandoski</i>	
Top-k Spatial Joins of Probabilistic Objects	566
<i>Vebjorn Ljosa, Ambuj K. Singh</i>	
The Cost of Serializability on Platforms That Use Snapshot Isolation	576
<i>Mohammad Alomari, Michael Cahill, Alan Fekete, Uwe Rohm</i>	
Network-Aware Join Processing in Global-Scale Database Federations	586
<i>Xiaodan Wang, Randal Burns, Andreas Terzis, Amol Deshpande</i>	

Table of Contents

Robust Runtime Optimization of Data Transfer in Queries Over Web Services	596
<i>Anastasios Gounaris, Christos Yfoulis, Rizos Sakellariou, Marios D. Dikaiakos</i>	
XML Processing in DHT Networks.....	606
<i>Serge Abiteboul, Ioana Manolescu, Neoklis Polyzotis, Nicoleta Preda, Chong Sun</i>	
Efficient Processing of XML Update Streams.....	616
<i>Leonidas Fegaras</i>	
XML Prefiltering as a String Matching Problem	626
<i>Christoph Koch, Stefanie Scherzinger, Michael Schmidt</i>	
Optimizing SQL Queries Over Text Databases	636
<i>Alpa Jain, AnHai Doan, Luis Gravano</i>	
Building Community Wikipedias: A Machine-Human Partnership Approach.....	646
<i>Pedro DeRose, Xiaoyong Chai, Byron J. Gao, Warren Shen, AnHai Doan, Philip Bohannon, Xiaojin Zhu</i>	
Keyword Search on Spatial Databases	656
<i>Ian De Felipe, Vagelis Hristidis, Naphtali Rishé</i>	
Simultaneous Equation Systems for Query Processing on Continuous-Time Data Streams.....	666
<i>Yanif Ahmad, Olga Papaemmanouil, Ugur Cetintemel, Jennie Rogers</i>	
Runtime Semantic Query Optimization for Event Stream Processing	676
<i>Luping Ding, Songting Chen, Elke A. Rundensteiner, Junichi Tatemura, Wang-Pin Hsiung, K. Selcuk Candan</i>	
Efficient Online Subsequence Searching in Data Streams Under Dynamic Time Warping Distance	686
<i>Mi Zhou, Man Hon Wong</i>	
Efficient Data Authentication in an Environment of Untrusted Third-Party Distributors	696
<i>Mikhail J. Atallah, YounSun Cho, Ashish Kundu</i>	
OptRR: Optimizing Randomized Response Schemes for Privacy-Preserving Data Mining	705
<i>Zhengli Huang, Wenliang Du</i>	
On the Anonymization of Sparse High-Dimensional Data.....	715
<i>Gabriel Ghinita, Yufei Tao, Panos Kalnis</i>	
On Anti-Corruption Privacy Preserving Publication.....	725
<i>Yufei Tao, Xiaokui Xiao, Jiexing Li, Donghui Zhang</i>	
Similar Document Detection with Limited Information Disclosure.....	735
<i>Wei Jiang, Mummoorthy Murugesan, Chris Clifton, Luo Si</i>	
k-Anonymization Revisited.....	744
<i>Aristides Gionis, Arnon Mazza, Tamir Tassa</i>	
Flow Algorithms for Parallel Query Optimization	754
<i>Amol Deshpande, Lisa Hellerstein</i>	
Self-Join Size Estimation in Large-Scale Distributed Data Systems	764
<i>Theoni Pitoura, Peter Triantafillou</i>	
Sideways Information Passing for Push-Style Query Processing.....	774
<i>Zachary G. Ives, Nicholas E. Taylor</i>	
COLR-Tree: Communication-Efficient Spatio-Temporal Indexing for a Sensor Data Web Portal.....	784
<i>Yanif Ahmad, Suman Nath</i>	
Region Sampling: Continuous Adaptive Sampling on Sensor Networks.....	794
<i>Song Lin, Benjamin Arai, Dimitrios Gunopulos, Gautam Das</i>	
Fast and Highly-Available Stream Processing Over Wide Area Networks.....	804
<i>Jeong-Hyon Hwang, Ugur Cetintemel, Stan Zdonik</i>	

Table of Contents

Approximate Joins for Data-Centric XML	814
<i>Nikolaus Augsten, Michael Bohlen, Curtis Dyreson, Johann Gamper</i>	
The Space Complexity of Processing XML Twig Queries Over Indexed Documents	824
<i>Mirit Shalem, Ziv Bar-Yossef</i>	
XML Index Recommendation with Tight Optimizer Coupling.....	833
<i>Iman Elghandour, Ashraf Aboulhaga, Daniel C. Zilio, Fei Chiang, Andrey Balmin, Kevin Beyer, Calisto Zuzarte</i>	
A Motion-Aware Approach to Continuous Retrieval of 3D Objects.....	843
<i>Mohammed Eunus Ali, Rui Zhang, Egemen Tanin, Lars Kulik</i>	
Continuous Content-Based Copy Detection Over Streaming Videos.....	853
<i>Ying Yan, Beng Chin Ooi, Aoying Zhou</i>	
Continuous Intersection Joins Over Moving Objects.....	863
<i>Rui Zhang, Dan Lin, Kotagiri Ramamohanarao, Elisa Bertino</i>	
Multiple Materialized View Selection for XPath Query Rewriting.....	873
<i>Nan Tang, Jeffrey Xu Yu, M. Tamer Ozsu, Byron Choi, Kam-Fai Wong</i>	
Explaining and Reformulating Authority Flow Queries.....	883
<i>Ramakrishna Varadarajan, Vagelis Hristidis, Louiqa Raschid</i>	
An Efficient Algorithm for Answering Graph Reachability Queries.....	893
<i>Yangjun Chen, Yibin Chen</i>	
Summarizing Graph Patterns.....	903
<i>Yong Liu, Jianzhong Li, Hong Gao</i>	
Fast Graph Pattern Matching	913
<i>Jiefeng Cheng, Jeffrey Xu Yu, Bolin Ding, Philip S. Yu, Haixun Wang</i>	
Stop Chasing Trends: Discovering High Order Models in Evolving Data	923
<i>Shixi Chen, Haixun Wang, Shuigeng Zhou, Philip S. Yu</i>	
An Algebraic Approach to Rule-Based Information Extraction	933
<i>Frederick Reiss, Sriram Raghavan, Rajasekar Krishnamurthy, Huaiyu Zhu, Shivakumar Vaithyanathan</i>	
Efficient Information Extraction Over Evolving Text Data	943
<i>Fei Chen, AnHai Doan, Jun Yang, Raghu Ramakrishnan</i>	
NAGA: Searching and Ranking Knowledge	953
<i>Gjergji Kasneci, Fabian M. Suchanek, Georgiana Ifrim, Maya Ramanath, Gerhard Weikum</i>	
TALE: A Tool for Approximate Large Graph Matching	963
<i>Yuanyuan Tian, Jignesh M. Patel</i>	
Probabilistic Verifiers: Evaluating Constrained Nearest-Neighbor Queries Over Uncertain Data.....	973
<i>Reynold Cheng, Jinchuan Chen, Mohamed F. Mokbel, Chi-Yin Chow</i>	
Fast and Simple Relational Processing of Uncertain Data.....	983
<i>Lyublena Antova, Thomas Jansen, Christoph Koch, Dan Olteanu</i>	
SHAROE: A Data Sharing Platform for Outsourced Enterprise Storage Environments	993
<i>Aameek Singh, Ling Liu</i>	
Accelerating Lookups in P2P Systems Using Peer Caching.....	1003
<i>Supratim Deb, Prakash Linga, Rajeev Rastogi, Anand Srinivasan</i>	
Diagnosing Estimation Errors in Page Counts Using Execution Feedback.....	1013
<i>Surajit Chaudhuri, Vivek Narasayya, Ravishankar Ramamurthy</i>	
Exploiting Lineage for Confidence Computation in Uncertain and Probabilistic Databases.....	1023
<i>Anish Das Sarma, Martin Theobald, Jennifer Widom</i>	

Table of Contents

Greedy List Intersection.....	1033
<i>Robert Krauthgamer, Aranyak Mehta, Vijayshankar Raman, Atri Rudra</i>	
Handling Non-Linear Polynomial Queries Over Dynamic Data	1043
<i>Shetal Shah, Krithi Ramamritham</i>	
Database Support for Probabilistic Attributes and Tuples.....	1053
<i>Sarvjeet Singh, Chris Mayfield, Rahul Shah, Sunil Prabhakar, Susanne Hambrusch, Jennifer Neville, Reynold Cheng</i>	
Scalable Rule-Based Gene Expression Data Classification	1062
<i>Mark A. Iwen, Willis Lang, Jignesh M. Patel</i>	
Querying and Managing Provenance Through User Views in Scientific Workflows.....	1072
<i>Olivier Biton, Sarah Cohen-Boulakia, Susan B. Davidson, Carmem S. Hara</i>	
Spatial Outsourcing for Location-Based Services.....	1082
<i>Yin Yang, Stavros Papadopoulos, Dimitris Papadias, George Kollios</i>	
P-Cube: Answering Preference Queries in Multi-Dimensional Space	1092
<i>Dong Xin, Jiawei Han</i>	
Efficient Rewriting Algorithms for Preference Queries.....	1101
<i>Periklis Georgiadis, Ioannis Kapantaidakis, Vassilis Christophides, Elhadji Mamadou Nguer, Nicolas Spyrtos</i>	
A Fast Similarity Join Algorithm Using Graphics Processing Units.....	1111
<i>Michael D. Lieberman, Jagan Sankaranarayanan, Hanan Samet</i>	
MobiQual: QoS-Aware Load Shedding in Mobile CQ Systems	1121
<i>Bugra Gedik, Kun-Lung Wu, Philip S. Yu, Ling Liu</i>	
Approximate Clustering on Distributed Data Streams	1131
<i>Qi Zhang, Jinze Liu, Wei Wang</i>	
Handling Uncertain Data in Array Database Systems	1140
<i>Tingjian Ge, Stan Zdonik</i>	
Just-In-Time Processing of Continuous Queries.....	1150
<i>Yin Yang, Dimitris Papadias</i>	
Online Filtering, Smoothing and Probabilistic Modeling of Streaming Data	1160
<i>Bhargav Kanagal, Amol Deshpande</i>	
Butterfly: Protecting Output Privacy in Stream Mining	1170
<i>Ting Wang, Ling Liu</i>	
XStream: A Signal-Oriented Data Stream Management System.....	1180
<i>Lewis Girod, Yuan Mei, Ryan Newton, Stanislav Rost, Arvind Thiagarajan, Hari Balakrishnan, Samuel Madden</i>	
Optimizing Star Join Queries for Data Warehousing in Microsoft SQL Server	1190
<i>Cesar A. Galindo-Legaria, Torsten Grabs, Sreenivas Gukal, Steve Herbert, Aleksandras Surna, Shirley Wang, Wei Yu, Peter Zabback, Shin Zhang</i>	
Integration of Server, Storage and Database Stack: Moving Processing Towards Data	1200
<i>Lin Qiao, Vijayshankar Raman, Inderpal Narang, Prashant Pandey, David Chambliss, Gene Fuh, James Ruddy, Ying-Lin Chen, Kou-Horng Yang, Fen-Ling Lin</i>	
Accurate and Efficient Inter-Transaction Dependency Tracking.....	1209
<i>Tzi-cker Chiueh, Shweta Bajpai</i>	
Native Web Browser Enabled SVG-Based Collaborative Multimedia Annotation for Medical Images	1219
<i>Fusheng Wang, Cornelius Rabsch, Peiya Liu</i>	
XQSE: An XQuery Scripting Extension for the AquaLogic Data Services Platform.....	1229
<i>Vinayak Borkar, Michael Carey, Daniel Engovatov, Dmitry Lychagin, Till Westmann, Warren Wong</i>	

Table of Contents

Implementing an Inference Engine for RDFS/OWL Constructs and User-Defined Rules in Oracle	1239
<i>Zhe Wu, George Eadon, Souripriya Das, Eugene Inseok Chong, Vladimir Kolovski, Melliyal Annamalai, Jagannathan Srinivasan</i>	
Towards Declarative Queries on Adaptive Data Structures.....	1249
<i>Nicolas Bruno, Pablo Castro</i>	
Filtered Indices and Their Use in Flexible Schema Scenarios	1259
<i>Srini Acharya, Cesar A. Galindo-Legaria, Milind Joshi, Babu Krishnaswamy, Stefano Stefani, Pawel Terlecki</i>	
Maintaining Connectivity in Dynamic Multimodal Network Models.....	1267
<i>Petko Bakalov, Erik Hoel, Wee-Liang Heng, Vassilis J. Tsotras</i>	
An Enhanced Extract-Transform-Load System for Migrating Data in Telecom Billing.....	1277
<i>Himanshu Agrawal, Girish Chafle, Sunil Goyal, Sumit Mittal, Sougata Mukherjea</i>	
An Overview of SQL Support in Workflow Products.....	1287
<i>Marko Vrhovnik, Holger Schwarz, Sylvia Radeschutz, Bernhard Mitschang</i>	
A Scalable Scheme for Bulk Loading Large RDF Graphs into Oracle.....	1297
<i>Souripriya Das, Eugene Inseok Chong, Zhe Wu, Melliyal Annamalai, Jagannathan Srinivasan</i>	
Orchid: Integrating Schema Mapping and ETL	1307
<i>Stefan Dessloch, Mauricio A. Hernandez, Ryan Wisnesky, Ahmed Radwan, Jindan Zhou</i>	
Merging Hierarchies Using Object Placement.....	1317
<i>Kai Zhao, Robert Ikeda, Hector Garcia-Molina</i>	
Model Management Engine for Data Integration with Reverse-Engineering Support.....	1319
<i>Michael N. Gubanov, Philip A. Bernstein, Alexander Moshchuk</i>	
Extracting Loosely Structured Data Records Through Mining Strict Patterns	1322
<i>Yipu Wu, Jing Chen, Qing Li</i>	
Detection of Shape Anomalies: A Probabilistic Approach Using Hidden Markov Models.....	1325
<i>Zheng Liu, Jeffrey Xu Yu, Lei Chen, Di Wu</i>	
Monitoring Network Evolution Using MDL	1328
<i>Jure Ferlez, Christos Faloutsos, Jure Leskovec, Dunja Mladenic, Marko Grobelnik</i>	
Database Exploration Using Join Paths.....	1331
<i>Cecilia M. Procopiuc, Divesh Srivastava</i>	
What-If OLAP Queries with Changing Dimensions.....	1334
<i>L.V.S. Lakshmanan, A. Russakovsky, V. Sashikanth</i>	
A General Framework for Fast Co-Clustering on Large Datasets Using Matrix Decomposition.....	1337
<i>Feng Pan, Xiang Zhang, Wei Wang</i>	
MUSE: Multi-Represented Similarity Estimation	1340
<i>Hans-Peter Kriegel, Peter Kunath, Alexey Pryakhin, Matthias Schubert</i>	
FLAME: Shedding Light on Hidden Frequent Patterns in Sequence Datasets	1343
<i>Sandeep Tata, Jignesh M. Patel</i>	
Mining Search-Phrase Definitions from Item Descriptions	1346
<i>Hung V. Nguyen, Hasan Davulcu</i>	
Rights Protection of Trajectory Datasets.....	1349
<i>Claudio Lucchese, Michail Vlachos, Deepak Rajan, Philip S. Yu</i>	
Privacy Preserving Joins.....	1352
<i>Yaping Li, Minghua Chen</i>	
Link Privacy in Social Networks	1355
<i>Aleksandra Korolova, Rajeev Motwani, Shubha U. Nabar, Ying Xu</i>	

Table of Contents

Protecting Databases from Query Flood Attacks	1358
<i>Anna Cinzia Squicciarini, Ivan Paloscia, Elisa Bertino</i>	
Secure Delta-Publishing of XML Content.....	1361
<i>Mohamed Nabeel, Elisa Bertino</i>	
On the Optimal Selection of k in the k-Anonymity Problem	1364
<i>Rinku Dewri, Indrajit Ray, Indrakshi Ray, Darrell Whitley</i>	
Anonymizing Streaming Data for Privacy Protection	1367
<i>Jianzhong Li, Beng Chin Ooi, Weiping Wang</i>	
Privately Updating Suppression and Generalization Based k-Anonymous Databases	1370
<i>Alberto Trombetta, Wei Jiang, Elisa Bertino, Lorenzo Bossi</i>	
InstantDB: Enforcing Timely Degradation of Sensitive Data.....	1373
<i>Nicolas Anceaix, Luc Bouganim, Harold van Heerde, Philippe Pucheral, Peter M.G. Apers</i>	
CASTLE: A Delay-Constrained Scheme for k_s-Anonymizing Data Streams	1376
<i>Jianneng Cao, Barbara Carminati, Elena Ferrari, Kian Lee Tan</i>	
Exponentially Decayed Aggregates on Data Streams	1379
<i>Graham Cormode, Flip Korn, Srikanta Tirthapura</i>	
Efficient Aggregate Computation Over Data Streams	1382
<i>Kanthi Nagaraj, K.V.M. Naidu, Rajeev Rastogi, Scott Satkin</i>	
Frequency Estimation Over Sliding Windows	1385
<i>Linfeng Zhang, Yong Guan</i>	
Online Failure Forecast for Fault-Tolerant Data Stream Processing	1388
<i>Xiaohui Gu, Spiros Papadimitriou, Philip S. Yu, Shu-Ping Chang</i>	
On Supporting Kleene Closure Over Event Streams	1391
<i>Daniel Gyllstrom, Jagrati Agrawal, Yanlei Diao, Neil Immerman</i>	
MOOLAP: Towards Multi-Objective OLAP	1394
<i>Shyam Antony, Ping Wu, Divyakant Agrawal, Amr El Abbadi</i>	
On Skylining with Flexible Dominance Relation	1397
<i>Tian Xia, Donghui Zhang, Yufei Tao</i>	
Designing Random Sample Synopses with Outliers	1400
<i>Philipp Rosch, Rainer Gemulla, Wolfgang Lehner</i>	
Efficiently Answering Probabilistic Threshold Top-k Queries on Uncertain Data	1403
<i>Ming Hua, Jian Pei, Wenjie Zhang, Xuemin Lin</i>	
Efficient Processing of Top-k Queries in Uncertain Databases	1406
<i>Ke Yi, Feifei Li, George Kollios, Divesh Srivastava</i>	
Adjoined Dimension Column Clustering to Improve Data Warehouse Query Performance.....	1409
<i>Xuedong Chen, Patrick O'Neil, Elizabeth O'Neil</i>	
Adaptive Segmentation for Scientific Databases.....	1412
<i>Milena Ivanova, Martin L. Kersten, Niels Nes</i>	
QShuffler: Getting the Query Mix Right.....	1415
<i>Mumtaz Ahmad, Ashraf Abounaga, Shivnath Babu, Kamesh Munagala</i>	
Adaptive Bitmap Indexes for Space-Constrained Systems	1418
<i>Rishi Rakesh Sinha, Marianne Winslett, Kesheng Wu, Kurt Stockinger, Arie Shoshani</i>	
Skyline-Based Peer-to-Peer Top-k Query Processing	1421
<i>Akrivi Vlachou, Christos Doulkeridis, Kjetil Norvaag, Michalis Vazirgiannis</i>	

Table of Contents

Adaptive Approximate Similarity Searching Through Metric Social Networks	1424
<i>Jan Sedmidubsky, Stanislav Barton, Vlastislav Dohnal, Pavel Zezula</i>	
Fixed-Precision Approximate Continuous Aggregate Queries in Peer-to-Peer Databases	1427
<i>Farnoush Banaei-Kashani, Cyrus Shahabi</i>	
Faster Join Enumeration for Complex Queries	1430
<i>Guido Moerkotte, Thomas Neumann</i>	
PermJoin: An Efficient Algorithm for Producing Early Results in Multi-Join Query Plans	1433
<i>Justin J. Levandoski, Mohamed E. Khalefa, Mohamed F. Mokbel</i>	
Processing Event-Monitoring Queries in Sensor Networks	1436
<i>Vassilis Stoumpos, Antonios Deligiannakis, Yannis Kotidis, Alex Delis</i>	
An Architecture for Query Optimization in Sensor Networks	1439
<i>Ixent Galpin, Christian Y.A. Brenninkmeijer, Farhana Jabeen, Alvaro A.A. Fernandes, Norman W. Paton</i>	
Optimizing Complex Event Processing Over RFID Data Streams	1442
<i>Qun Chen, Zhanhuai Li, Hailong Liu</i>	
Efficient Data Interpretation and Compression Over RFID Streams	1445
<i>Richard Cocci, Thanh Tran, Yanlei Diao, Prashant Shenoy</i>	
Outlier-Aware Data Aggregation in Sensor Networks	1448
<i>Antonios Deligiannakis, Vassilis Stoumpos, Yannis Kotidis, Vasilis Vassalos, Alex Delis</i>	
Efficient Query Processing in Large Traffic Networks	1451
<i>Hans-Peter Kriegel, Peer Kroger, Peter Kunath, Matthias Renz, Tim Schmidt</i>	
Optimal-Nearest-Neighbor Queries	1454
<i>Yunjun Gao, Jing Zhang, Gencai Chen, Qing Li, Shen Liu, Chun Chen</i>	
Convoy Queries in Spatio-Temporal Databases	1457
<i>Hoyoung Jeung, Heng Tao Shen, Xiaofang Zhou</i>	
On High Dimensional Indexing of Uncertain Data	1460
<i>Charu C. Aggarwal, Philip S. Yu</i>	
Pattern Matching Over Cloaked Time Series	1462
<i>Xiang Lian, Lei Chen, Jeffrey Xu Yu</i>	
Satisfying Complex Data Needs using Pull-Based Online Monitoring of Volatile Data Sources.....	1465
<i>Haggai Roitman, Avigdor Gal, Louiqa Raschid</i>	
Processing Diagnosis Queries: A Principled and Scalable Approach	1468
<i>Shivnath Babu, Songyun Duan, Kamesh Munagala</i>	
Capturing Approximated Data Delivery Tradeoffs.....	1471
<i>Haggai Roitman, Avigdor Gal, Louiqa Raschid</i>	
Skippy: Enabling Long-Lived Snapshots of the Long-Lived Past.....	1474
<i>Ross Shaull, Liuba Shrira, Hao Xu</i>	
Adjourn State Concurrency Control Avoiding Time-Out Problems in Atomic Commit Protocols.....	1477
<i>Sebastian Obermeier, Stefan Bottcher, Martin Hett, Panos K. Chrysanthis, George Samaras</i>	
Probabilistic Event Extraction from RFID Data	1480
<i>Nodira Khoussainova, Magdalena Balazinska, Dan Suciu</i>	
Mobile Filter: Exploring Migration of Filters for Error-Bounded Data Collection in Sensor Networks.....	1483
<i>Dan Wang, Jianliang Xu, Jiangchuan Liu, Feng Wang</i>	
PageChaser: A Tool for the Automatic Correction of Broken Web Links.....	1486
<i>Atsuyuki Morishima, Akiyoshi Nakamizo, Toshinari Iida, Shigeo Sugimoto, Hiroyuki Kitagawa</i>	

Table of Contents

Index Design for Dynamic Personalized PageRank	1489
<i>Amit Pathak, Soumen Chakrabarti, Manish Gupta</i>	
Quality-Aware Retrieval of Data Objects from Autonomous Sources for Web-Based Repositories	1492
<i>Houtan Shirani-Mehr, Chen Li, Gang Liang, Michal Shmueli-Scheuer</i>	
Efficient Discovery of Authoritative Resources	1495
<i>Ravi Kumar, Kevin Lang, Cameron Marlow, Andrew Tomkins</i>	
Automatically Extracting Form Labels	1498
<i>Hoa Nguyen, Eun Yong Kang, Juliana Freire</i>	
Answering Keyword Queries on XML Using Materialized Views	1501
<i>Ziyang Liu, Yi Chen</i>	
An Inflationary Fixed Point Operator in XQuery	1504
<i>Loredana Afanasiev, Torsten Grust, Maarten Marx, Jan Rittinger, Jens Teubner</i>	
Grouping and Optimization of XPath Expressions in System RX	1507
<i>Andrey Balmin, Fatma Ozcan, Ashutosh Singh, Edison Ting</i>	
Hierarchical Indexing Approach to Support XPath Queries	1510
<i>Nan Tang, Jeffrey Xu Yu, M. Tamer Ozsu, Kam-Fai Wong</i>	
A Query Processing Architecture for an XML Data Warehouse	1513
<i>Nuwee Wiwatwattana, H.V. Jagadish</i>	
A Clustered Index Approach to Distributed XPath Processing	1516
<i>Georgia Koloniari, Evaggelia Pitoura</i>	
AxPRE Summaries: Exploring the (Semi-)Structure of XML Web Collections	1519
<i>Mariano P. Consens, Flavio Rizzolo, Alejandro A. Vaisman</i>	
Correlation-Based Attribute Outlier Detection in XML	1522
<i>Judice L.Y. Koh, Mong Li Lee, Wynne Hsu, Wee Tiong Ang</i>	
Ontology-Aware Search on XML-Based Electronic Medical Records	1525
<i>Fernando Farfan, Vagelis Hristidis, Anand Ranganathan, Redmond P. Burke</i>	
Query-Aware Partitioning for Monitoring Massive Network Data Streams	1528
<i>Theodore Johnson, S. Muthukrishnan, Vladislav Shkapenyuk, Oliver Spatscheck</i>	
Optimizing Hierarchical Access in OLAP Environment	1531
<i>Lipyeow Lim, Bishwaranjan Bhattacharjee</i>	
Automatic Result Verification for the Functional Testing of a Query Language	1534
<i>Carsten Binnig, Donald Kossmann, Eric Lo, Angel Saenz-Badillos</i>	
Improving Information Access for a Community of Practice Using Business Process as Context	1537
<i>Yu Deng, Murthy Devarakonda, Nithya Rajamani, Wlodek Zadrozny</i>	
DescribeX: Interacting with AxPRE Summaries	1540
<i>M.S. Ali, Mariano P. Consens, Shahan Khatchadourian, Flavio Rizzolo</i>	
COSTA: Adaptive Indexing for Terms in a Large-Scale Distributed System	1544
<i>Aoying Zhou, Rong Zhang, Quang Hieu Vu, Weining Qian</i>	
IMPRECISE: Good-is-Good-Enough Data Integration	1548
<i>Ander de Keijzer, Maurice van Keulen</i>	
SPARK: A Keyword Search Engine on Relational Databases	1552
<i>Yi Luo, Wei Wang, Xuemin Lin</i>	
Astoria: A Programming Model for Data on the Web	1556
<i>Pablo Castro, Anil Nori</i>	

Table of Contents

Environmental Tomography: Ubiquitous Sensing with Mobile Devices	1560
<i>Stacy Patterson, Bassam Bamieh, Amr El Abbadi</i>	
OptimAX: Efficient Support for Data-Intensive Mash-Ups	1564
<i>Serge Abiteboul, Ioana Manolescu, S. Zoupanos</i>	
Graphitti: An Annotation Management System for Heterogeneous Objects.....	1568
<i>Sandeep Gupta, Christopher Condit, Amarnath Gupta</i>	
Distributed Monitoring of Peer-to-Peer Systems.....	1572
<i>Serge Abiteboul, Bogdan Marinoiu, Pierre Bourhis</i>	
SMART: A System for Online Monitoring Large Volumes of Network Traffic.....	1576
<i>Aoying Zhou, Ying Yan, Xueqing Gong, Jianlong Chang, Dai Dai</i>	
Toward Simulation-Based Optimization in Data Stream Management Systems	1580
<i>Christoph Heinz, Jurgen Kramer, Tobias Riemenschneider, Bernhard Seeger</i>	
PGM/F: A Framework for the Optimization of Data Processing in Business Processes.....	1584
<i>Marko Vrhovnik, Oliver Suhre, Stephan Ewen, Holger Schwarz</i>	
Fuzzy Multi-Dimensional Search in the Wayfinder File System.....	1588
<i>Christopher Peery, Wei Wang, Amelie Marian, Thu D. Nguyen</i>	
MED: A Multimedia Event Database for 3D Crime Scene Representation and Analysis	1592
<i>Marcin Kwietniewski, Stephanie Wilson, Anna Topol, Sunbir Gill, Jarek Gryz, Michael Jenkin, Piotr Jasiobedzki, Ho-Kong Ng</i>	
DIPBench Toolsuite: A Framework for Benchmarking Integration Systems.....	1596
<i>Matthias Bohm, Dirk Habich, Wolfgang Lehner, Uwe Wloka</i>	
Managing Biological Data Using bdbms.....	1600
<i>Mohamed Y. Eltabakh, Mourad Ouzzani, Walid G. Aref, Ahmed K. Elmagarmid, Yasin Laura-Silva, Muhammad U. Arshad, David Salt, Ivan Baxter</i>	
Mining Gene Expression Database for Primary Human Disease Tissues	1604
<i>Andrew Campen, Yuni Xia, Dan Rigsby, Ying Guo, Xingdong Feng, Eric W. Su, Mathew Palakal, Shuyu Li</i>	
Mining Views: Database Views for Data Mining.....	1608
<i>Hendrik Blockeel, Toon Calders, Elisa Fromont, Bart Goethals, Adriana Prado</i>	
PictureBook: A Text-and-Image Summary System for Web Search Result.....	1612
<i>Mei Wang, Hongtao Xu, Guoyu Hao, Xiangdong Zhou, Wei Wang, Qi Zhang, Baile Shi</i>	
Enhanced Business Intelligence Using EROCS.....	1616
<i>M. Bhide, V. Chakravarthy, A. Gupta, H. Gupta, M. Mohania, K. Puniyani, P. Roy, S. Roy, V. Sengar</i>	
T-Time: Threshold-Based Data Mining on Time Series.....	1620
<i>Johannes Assfalg, Hans-Peter Kriegel, Peer Kroger, Peter Kunath, Alexey Pryakhin, Matthias Renz</i>	
RAD: A Scalable Framework for Annotator Development	1624
<i>Sanjeet Khaitan, Ganesh Ramakrishnan, Sachindra Joshi, Anup Chalamalla</i>	
SPOT: A System for Detecting Projected Outliers from High-Dimensional Data Streams.....	1628
<i>Ji Zhang, Qigang Gao, Hai Wang</i>	