



22nd International Conf. on Data Engineering
Atlanta, Georgia, USA, 2006



**Proceedings of the
22nd International Conference on
Data Engineering**

**April 3-7, 2006
Atlanta, Georgia, USA**

TABLE OF CONTENTS

RESEARCH SESSION 1: VIEWS

MiniCount: Efficient Rewriting of COUNT-Queries Using Views	1
<i>Václav Lín, Vasilis Vassalos, Prodromos Malakasiotis</i>	
Updates through Views: A New Hope	12
<i>Yannis Kotidis, Divesh Srivastava, Yannis Velegarakis</i>	
Learning from Aggregate Views	24
<i>Bee-Chung Chen, Lei Chen, Raghu Ramakrishnan, David R. Musicant</i>	

RESEARCH SESSION 2: DATA WAREHOUSE (1)

C-Cubing: Efficient Computation of Closed Cubes by Aggregation-Based Checking	34
<i>Dong Xin, Zheng Shao, Jiawei Han, Hongyan Liu</i>	
A Primitive Operator for Similarity Joins in Data Cleaning	44
<i>Surajit Chaudhuri, Venkatesh Ganti, Raghav Kaushik</i>	
Techniques for Warehousing of Sample Data	56
<i>Paul G. Brown, Peter J. Haas</i>	

RESEARCH SESSION 3: QUERY PROCESSING AND UNCERTAINTY REASONING

Working Models for Uncertain Data	68
<i>Anish Das Sarma, Omar Benjelloun, Alon Halevy, Jennifer Widom</i>	
Reasoning about Approximate Match Query Results	80
<i>Sudipto Guha, Nick Koudas, Divesh Srivastava, Xiaohui Yu</i>	
The Gauss-Tree: Efficient Object Identification in Databases of Probabilistic Feature Vectors	92
<i>Christian Böhm, Alexey Pryakhin, Matthias Schubert</i>	

RESEARCH SESSION 4: INDEXING AND OPTIMIZATION

Finding Fastest Paths on A Road Network with Speed Patterns	104
<i>Evangelos Kanoulas, Yang Du, Tian Xia, Donghui Zhang</i>	
Approximation Techniques for Indexing the Earth Mover™'s Distance in Multimedia Databases	114
<i>Ira Assent, Andrea Wenning, Thomas Seidl</i>	
Indexing for Dynamic Abstract Regions	126
<i>Joxan Jaffar, Roland H. C. Yap, Kenny Q. Zhu</i>	

RESEARCH SESSION 5: XML AND SEMI-STRUCTURED DATA

Efficient Processing of Updates in Dynamic XML Data	137
<i>Changqing Li, Tok Wang Ling, Min Hu</i>	
A Complete and Efficient Algebraic Compiler for XQuery	147
<i>Christopher Ré, Jérôme Siméon, Mary Fernández</i>	
Making Designer Schemas with Colors	159
<i>Nuwee Wiwatwattana, H.V. Jagadish, Laks V. S. Lakshmanan, Divesh Srivastava</i>	

RESEARCH SESSION 6: DATA MINING AND OPTIMIZATION

Mining Actionable Patterns by Role Models	170
<i>Ke Wang, Yuelong Jiang, Alexander Tuzhilin</i>	
A Systematic Approach for Optimizing Complex Mining Tasks on Multiple Databases	180
<i>Ruoming Jin, Gagan Agrawal</i>	
New Sampling-Based Estimators for OLAP Queries	190
<i>Ruoming Jin, Leo Glimcher, Chris Jermaine, Gagan Agrawal</i>	
Efficiently Evaluating Order Preserving Similarity Queries over Historical Market-Basket Data	200
<i>Reza Sherkat, Davood Rafiei</i>	

RESEARCH SESSION 7: QUERY PROCESSING AND QUERY OPTIMIZATION

End-biased Samples for Join Cardinality Estimation	212
<i>Cristian Estan, Jeffrey F. Naughton</i>	
Laws for Rewriting Queries Containing Division Operators	224
<i>Ralf Rantzau, Christoph Mangold</i>	
R-trees with Update Memos	234
<i>Xiaopeng Xiong, Walid G. Aref</i>	
Compiled Query Execution Engine using JVM	244
<i>Jun Rao, Hamid Pirahesh, C. Mohan, Guy Lohman</i>	

RESEARCH SESSION 8: DATA PRIVACY AND SECURITY

I-Diversity: Privacy beyond K-Anonymity	256
<i>Ashwin Machanavajjhala, Johannes Gehrke, Daniel Kifer, Muthuramakrishnan Venkitasubramaniam</i>	
Mondrian Multidimensional K-Anonymity	268
<i>Kristen LeFevre, David J. DeWitt, Raghu Ramakrishnan</i>	
Sovereign Joins	279
<i>Rakesh Agrawal, Dmitri Asonov, Murat Kantarcioglu, Yaping Li</i>	
Privacy Preserving Query Processing using Third Parties	291
<i>Fatih Emekci, Divyakant Agrawal, Amr El Abbadi, Aziz Gülbiden</i>	

RESEARCH SESSION 9: DATA INTEGRATION AND DATABASE INTEROPERABILITY

Efficient Batch Top-k Search for Dictionary-Based Entity Recognition	301
<i>Amit Chandel, P. C. Nagesh, Sunita Sarawagi</i>	
Integrating Unstructured Data into Relational Databases	311
<i>Imran R. Mansuri, Sunita Sarawagi</i>	
Clean Answers over Dirty Databases: A Probabilistic Approach	323
<i>Periklis Andritsos, Ariel Fuxman, Renée J. Miller</i>	

RESEARCH SESSION 10: WEB SERVICES AND APPLICATIONS

A Syntactic Rule Based Approach to Web Service Composition	335
<i>Ken Pu, Vagelis Hristidis, Nick Koudas</i>	
Hilda: A High-Level Language for Data-Driven Web Applications	345
<i>Fan Yang, Jayavel Shanmugasundaram, Mirek Riedewald, Johannes Gehrke, Alan Demers</i>	
UNIT: User-centric Transaction Management in Web-Database Systems	357
<i>Huiming Qu, Alexandros Labrinidis, Daniel Mossé</i>	

RESEARCH SESSION 11: TEMPORAL AND SPATIAL DATA MANAGEMENT (1)

VBI-Tree: A Peer-to-Peer Framework for Supporting Multi-Dimensional Indexing Schemes	367
<i>H.V. Jagadish, Beng Chin Ooi, Quang Hieu Vu, Rong Zhang, Aoying Zhou</i>	
Transaction Time Support inside a Database Engine	377
<i>David Lomet, Roger Barga, Mohamed F. Mokbel, German Shegalov, Rui Wang, Yunyue Zhu</i>	
Expiration Times for Data Management	389
<i>Albrecht Schmidt, Christian S. Jensen, Simonas Saltenis</i>	

RESEARCH SESSION 12: QUERY OPTIMIZATION AND DATA STRUCTURES

Scalable Exploration of Physical Database Design	400
<i>Arnd Christian König, Shubha U. Nabar</i>	
Closure-Tree: An Index Structure for Graph Queries	411
<i>Huahai He, Ambuj K. Singh</i>	
ISOMER: Consistent Histogram Construction Using Query Feedback	423
<i>Utkarsh Srivastava, Peter J. Haas, Volker Markl, Nimrod Megiddo, Marsel Kutsch, Tam Minh Tran</i>	

RESEARCH SESSION 13: DISTRIBUTED AND PEER TO PEER DATA MANAGEMENT

Counting at Large: Efficient Cardinality Estimation in Internet-Scale Data Networks	435
<i>Nikos Ntarmos, Peter Triantafillou, Gerhard Weikum</i>	
Probabilistic Message Passing in Peer Data Management Systems	445
<i>Philippe Cudré-Mauroux, Karl Aberer, Andras Feher</i>	

Approximating Aggregation Queries in Peer-to-Peer Networks	457
<i>Benjamin Arai, Gautam Das, Dimitrios Gunopulos, Vana Kalogeraki</i>	
Distributed Evaluation of Continuous Equi-join Queries over Large Structured Overlay Networks	469
<i>Stratos Idreos, Christos Tryfonopoulos, Manolis Koubarakis</i>	

RESEARCH SESSION 14: WEB QUERIES

WebIQ: Learning from the Web to Match Deep-Web	481
<i>Wensheng Wu, AnHai Doan, Clement Yu</i>	
Answering Imprecise Queries over Autonomous Web Databases	491
<i>Ullas Nambiar, Subbarao Kambhampati</i>	
Merging Source Query Interfaces on Web Databases	501
<i>Eduard Dragut, Wensheng Wu, Prasad Sistla, Clement Yu, Weiyi Meng</i>	
Query Selection Techniques for Efficient Crawling of Structured Web Sources.....	511
<i>Ping Wu, Ji-Rong Wen, Huan Liu, Wei-Ying Ma</i>	

RESEARCH SESSION 15: STREAM PROCESSING (1)

Approximate Data Collection in Sensor Networks using Probabilistic Models.....	521
<i>David Chu, Amol Deshpande, Joseph M. Hellerstein, Wei Hong</i>	
Network-Aware Operator Placement for Stream-Processing Systems	533
<i>Peter Pietzuch, Jonathan Ledlie, Jeffrey Shneidman, Mema Roussopoulos, Matt Welsh, Margo Seltzer</i>	
Unifying the Processing of XML Streams and Relational Data Streams	545
<i>Xin Zhou, Hetal Thakkar, Carlo Zaniolo</i>	
Space-efficient Relative Error Order Sketch over Data Streams	557
<i>Ying Zhang, Xuemin Lin, Jian Xu, Flip Korn, Wei Wang</i>	

RESEARCH SESSION 16: XML AND XPATH

Designing and Evaluating an XPath Dialect for Linguistic Queries.....	569
<i>Steven Bird, Yi Chen, Susan B. Davidson, Haejoong Lee, Yifeng Zheng</i>	
Predicate-based Filtering of XPath Expressions	579
<i>Shuang Hou, H.-Arno Jacobsen</i>	
An Estimation System for XPath Expressions.....	589
<i>Hanyu Li, Mong Li Lee, Wynne Hsu, Gao Cong</i>	

RESEARCH SESSION 17: STREAM PROCESSING (2)

Adaptive Reorganization of Coherency-Preserving Dissemination Tree for Streaming Data	600
<i>Yongluan Zhou, Beng Chin Ooi, Kian-Lee Tan, Feng Yu</i>	
Declarative Network Monitoring with an Underprovisioned Query Processor	610
<i>Frederick Reiss, Joseph M. Hellerstein</i>	

What's Different: Distributed, Continuous Monitoring of Duplicate-Resilient Aggregates on Data Streams	622
<i>Graham Cormode, S. Muthukrishnan, Wei Zhuang</i>	

RESEARCH SESSION 18: DATABASE SYSTEM INTERNALS AND PERFORMANCE

Extending RDBMSs to Support Sparse Datasets Using an Interpreted Attribute Storage Format.....	632
<i>Jennifer L. Beckmann, Alan Halverson, Rajasekar Krishnamurthy, Jeffrey F. Naughton</i>	
Super-Scalar RAM-CPU Cache Compression	642
<i>Marcin Zukowski, Sándor Héman, Niels Nes, Peter Boncz</i>	
How to Determine a Good Multi-Programming Level for External Scheduling.....	654
<i>Bianca Schroeder, Mor Harchol-Balter, Arun Iyengar, Erich Nahum, Adam Wierman</i>	

RESEARCH SESSION 19: XML DATABASES AND XML QUERY OPTIMIZATION

XSeed: Accurate and Fast Cardinality Estimation for XPath Queries	666
<i>Ning Zhang, Tamer Ozsu, Ashraf Aboulnaga, Ihab F. Ilyas</i>	
Estimating XML Structural Join Size Quickly and Economically	678
<i>Cheng Luo, Zhewei Jiang, Wen-Chi Hou, Feng Yan, Chih-Fang Wang</i>	
XCluster Synopses for Structured XML Content	688
<i>Neoklis Polyzotis, Minos Garofalakis</i>	

RESEARCH SESSION 20: SKYLINE QUERY PROCESSING

Robust Cardinality and Cost Estimation for Skyline Operator.....	700
<i>Surajit Chaudhuri, Nilesh Dalvi, Raghav Kaushik</i>	
SUBSKY: Efficient Computation of Skylines in Subspaces	710
<i>Yufei Tao, Xiaokui Xiao, Jian Pei</i>	
Skyline Queries Against Mobile Lightweight Devices in MANETs.....	720
<i>Zhiyong Huang, Christian S. Jensen, Hua Lu, Beng Chin Ooi</i>	

RESEARCH SESSION 21: QUERY MANAGEMENT

Approximately Processing Multi-granularity Aggregate Queries over a Data Stream	731
<i>Shouke Qin, Weining Qian, Aoying Zhou</i>	
A Sampling-Based Approach to Optimizing Top-k Queries in Sensor Networks	741
<i>Adam Silberstein, Rebecca Braynard, Carla Ellis, Kamesh Munagala, Jun Yang</i>	
Précis: The Essence of a Query Answer	752
<i>Georgia Koutrika, Alkis Simitsis, Yannis Ioannidis</i>	

RESEARCH SESSION 22: UBIQUITOUS DATA MANAGEMENT

Processing Window Queries in Wireless Sensor Networks	762
<i>Yingqi Xu, Wang-Chien Lee, Jianliang Xu, Gail Mitchell</i>	
Effective Density Queries of Continuously Moving Objects.....	772
<i>Christian S. Jensen, Dan Lin, Beng Chin Ooi, Rui Zhang</i>	
Efficient Aggregation of Ranked Inputs	783
<i>Nikos Mamoulis, Kit Hung Cheng, Man Lung Yiu, David W. Cheung</i>	

RESEARCH SESSION 23: GRAPH DATABASES AND ALGORITHMS

CLAN: An Algorithm for Mining Closed Cliques from Large Dense Graph Databases.....	795
<i>Jianyong Wang, Zhiping Zeng, Lizhu Zhou</i>	
A Partition-Based Approach to Graph Mining	805
<i>Junmei Wang, Wynne Hsu, Mong Li Lee, Chang Sheng</i>	
Dual Labeling: Answering Graph Reachability Queries in Constant Time	815
<i>Haixun Wang, Hao He, Jun Yang, Philip S. Yu, Jeffrey Xu Yu</i>	

RESEARCH SESSION 24: NEAREST NEIGHBOR SEARCH IN TEMPORAL AND SPATIAL DATABASES

Reverse Nearest Neighbors Search in Ad-hoc Subspaces	827
<i>Man Lung Yiu, Nikos Mamoulis</i>	
Continuous Reverse Nearest Neighbor Monitoring	837
<i>Tian Xia, Donghui Zhang</i>	
Surface K-NN Query Processing.....	847
<i>Ke Deng, Xiaofang Zhou, Heng Tao Shen, Kai Xu, Xuemin Lin</i>	

RESEARCH SESSION 25: STREAM QUERY OPTIMIZATION

An Efficient XPath Query Processor for XML Streams.....	858
<i>Yi Chen, Susan B. Davidson, Yifeng Zheng</i>	
SketchTree: Approximate Tree Pattern Counts over Streaming Labeled Trees	870
<i>Praveen Rao, Bongki Moon</i>	
Characterizing and Exploiting Reference Locality in Data Stream Applications.....	882
<i>Feifei Li, Ching Chang, George Kollios, Azer Bestavros</i>	

RESEARCH SESSION 26: ADVANCED QUERY PROCESSING

MONDRIAN: Annotating and Querying Databases through Colors and Blocks	894
<i>Floris Geerts, Anastasios Kementsietsidis, Diego Milano</i>	
Warehousing and Analyzing Massive RFID Data Sets	904
<i>Hector Gonzalez, Jiawei Han, Xiaolei Li, Diego Klabjan</i>	

RESEARCH SESSION 27: TEMPORAL AND SPATIAL DATA MANAGEMENT (2)

Query Decomposition: A Multiple Neighborhood Approach to Relevance Feedback Processing in Content-Based Image Retrieval	914
<i>Kien A. Hua, Ning Yu, Danzhou Liu</i>	
Nearest Surrounding Queries	925
<i>Ken C. K. Lee, Wang-Chien Lee, Hong Va Leong</i>	
Closest-Point-of-Approach Join for Moving Object Histories	935
<i>Subramanian Arumugam, Christopher Jermaine</i>	

RESEARCH SESSION 28: SCIENTIFIC AND BIOLOGICAL DATABASES AND BIOINFORMATICS

Declarative Querying for Biological Sequences	945
<i>Sandeep Tata, Jignesh M. Patel, James S. Friedman, Anand Swaroop</i>	
Searching Substructures with Superimposed Distance	957
<i>Xifeng Yan, Feida Zhu, Jiawei Han, Philip S. Yu</i>	
Mining Shifting-and-Scaling Co-Regulation Patterns on Gene Expression Profiles	967
<i>Xin Xu, Ying Lu, Anthony K.H. Tung, Wei Wang</i>	

INDUSTRY SESSION 1: DATABASES, SERVICES, AND PROCESSES

AutoGlobe: An Automatic Administration Concept for Service-Oriented Database Applications	977
<i>Stefan Seltzsa, Daniel Gmach, Stefan Krompass, Alfons Kemper</i>	
Collaborative Business Process Support in IHE XDS through ebXML Business Process	989
<i>Asuman Dogac, Veli Bicer, Alper Okcan</i>	
Taming Compliance with Sarbanes-Oxley Internal Controls Using Database Technology	999
<i>Rakesh Agrawal, Christopher Johnson, Jerry Kiernan, Frank Leymann</i>	

INDUSTRY SESSION 2: RDF, ONTOLOGIES, METADATA

RDF Object Type and Reification in the Database	1009
<i>Nicole Alexander, Siva Ravada</i>	
RDF/RDFS-based Relational Database Integration	1019
<i>Huajun Chen, Zhaohui Wu, Heng Wang, Yuxin Mao</i>	
Supporting Keyword Columns with Ontology-based Referential Constraints in DBMS	1029
<i>Eugene Chong, Souripriya Das, George Eadon, Jagannathan Srinivasan</i>	
Experiment Management with Metadata-based Integration for Collaborative Scientific Research	1039
<i>Fusheng Wang, Peiya Liu, John Pearson, Fred Azar, Gerald Madlmayr</i>	

INDUSTRY SESSION 3: QUERYING AND LOGGING

SQL to XQuery Translation in the AquaLogic Data Services Platform	1050
<i>Sunil Jigyasu, Sujeet Banerjee, Vinayak Borkar, Michael Carey, Kanad Dixit, Anil Malkani, Sachin Thatte</i>	
On Pushing Multilingual Query Operators into Relational Engines.....	1062
<i>A Kumaran, Pavan K. Chowdary, Jayant R. Haritsa</i>	
SIREN: A Memory-Conserving, Snapshot-Consistent Checkpoint Algorithm for in-Memory Databases.....	1072
<i>Antti-Pekka Lienes, Antoni Wolski</i>	

INDUSTRY SESSION 4: DATA WAREHOUSING AND ANALYSIS

Space-Partitioning Trees in PostgreSQL: Realization and Performance	1084
<i>M.Y. Eitabakh, Ramy Eltarras, Walid G. Aref</i>	
Automatic Sales Lead Generation from Web Data.....	1096
<i>Ganesh Ramakrishnan, Sachindra Joshi, Sumit Negi, Raghu Krishnapuram, Sreeram Balakrishnan</i>	
Load Balancing for Multi-tiered Database Systems through Autonomic Placement of Materialized Views.....	1106
<i>Wen-Syan Li, Daniel C. Zilio, Vishal Batra, Mahadevan Subramanian, Calisto Zuzarte, Inderpal Narang</i>	

ADVANCED SEMINAR 1:

Schema and Data Translation	1118
<i>Paolo Atzeni</i>	

ADVANCED SEMINAR 2:

Foundations of Automated Database Tuning	1119
<i>Surajit Chaudhuri, Gerhard Weikum</i>	

ADVANCED SEMINAR 3:

Models and Methods for Privacy-Preserving Data Publishing and Analysis	1120
<i>Johannes Gehrke</i>	

ADVANCED SEMINAR 4:

Mining, Indexing, and Similarity Search in Graphs and Complex Structures.....	1121
<i>Jiawei Han, Xifeng Yan, Philip S. Yu</i>	

SESSION SEMINAR 5:

Query co-Processing on Commodity Hardware	1122
<i>Anastassia Ailamaki, Naga K. Govindaraju, Dinesh Manocha</i>	

POSTER SESSION 1: DATA INTEGRATION, DATA MINING, DATA WAREHOUSING

Efficient Continuous Skyline Computation	1123
<i>Michael Morse, Jignesh M. Patel, William I. Grosky</i>	
Detecting Duplicates in Complex XML Data.....	1126
<i>Melanie Weis, Felix Naumann</i>	
SaveRF: Towards Efficient Relevance Feedback Search	1129
<i>Heng Tao Shen, Beng Chin Ooi, Kian-Lee Tan</i>	
On the Inverse Classification Problem and its Applications	1132
<i>Charu C. Aggarwal, Chen Chen, Jiawei Han</i>	
MIC Framework: An Information-Theoretic Approach to Quantitative Association Rule Mining	1135
<i>Yiping Ke, James Cheng, Wilfred Ng</i>	
Efficient Discovery of Emerging Frequent Patterns in Arbitrary Windows on Data Streams.....	1138
<i>Xiaoming Jin, Xinqiang Zuo, Kwok-Yan Lam, Jianmin Wang, Jianguang Sun</i>	
Top-down Mining of Interesting Patterns from Very High Dimensional Data	1142
<i>Hongyan Liu, Jiawei Han, Dong Xin, Zheng Shao</i>	
Mining Dense Periodic Patterns in Time Series Data	1145
<i>Chang Sheng, Wynne Hsu, Mong Li Lee</i>	

POSTER SESSION 2: DATA PRIVACY AND SECURITY

Private Updates to Anonymous Databases.....	1148
<i>Alberto Trombetta, Elisa Bertino</i>	
Provable Security for Outsourcing Database Operations	1151
<i>Sergei Evdokimov, Matthias Fischmann, Oliver Günther</i>	
Composition and Disclosure of Unlinkable Distributed Databases	1155
<i>Bradley Malin, Latanya Sweeney</i>	
A Technique for Optimal Adaptation of Time-Dependant Workflow with Security.....	1158
<i>Basit Shafiq, Arjmand Samuel, Elisa Bertino, Arif Ghafoor</i>	

POSTER SESSION 3: WEB DATA MANAGEMENT AND WORKFLOW

Segmentation of Publication Records of Authors from the Web	1161
<i>Wei Zhang, Neil Smalheiser, Vette Torvik, Clement Yu</i>	
Text Classification Improved through Automatically Extracted Sequences	1164
<i>Dou Shen, Jian-Tao Sun, Qiang Yang, Hui Zhao, Zheng Chen</i>	

Holistic Query Interface Matching using Parallel Schema Matching	1167
<i>Weifeng Su, Jiyang Wang, Frederick Lochovsky</i>	
Extracting Objects from the Web.....	1171
<i>Zaiqing Nie, Fei Wu, Ji-Rong Wen, Wei-Ying Ma</i>	
On the Controlled Evolution of Process Choreographies.....	1174
<i>Stefanie Rinderle, Andreas Wombacher, Manfred Reichert</i>	
Automating the Design and Construction of Query Forms.....	1177
<i>Magesh Jayapandian, H.V. Jagadish</i>	

POSTER SESSION 4: XML DATA MANAGEMENT

U-Filter: A Lightweight XML View Update Checker.....	1180
<i>Ling Wang, Elke A. Rudensteiner, and Murali Mani</i>	
XCut: Indexing XML Data for Efficient Twig Evaluation.....	1183
<i>Simon Sheu, Nigel Wu</i>	
Algebraic Optimization of Nested XPath Expressions.....	1187
<i>Matthias Brantner, Sven Helmer, Carl-Christian Kanne, Guido Moerkotte</i>	
Incremental Maintenance of Materialized XQuery Views	1190
<i>Maged El-Sayed, Elke A. Rudensteiner, Murali Mani</i>	
Don't be a Pessimist: Use Snapshot based Concurrency Control for XML.....	1193
<i>Zeeshan Sardar, Bettina Kemme</i>	
Using XML to Build Efficient Transaction-Time Temporal Database Systems on Relational Databases.....	1196
<i>Fusheng Wang, Xin Zhou, Carlo Zaniolo</i>	

POSTER SESSION 5: PEER TO PEER AND MOBILE DATA MANAGEMENT

Partial Selection Query in Peer-to-Peer Databases	1200
<i>Farnoush Banaei-Kashani, Cyrus Shahabi</i>	
HiWaRPP --- Hierarchical Wavelet-based Retrieval on Peer-to-Peer Network.....	1203
<i>Mihai Lupu, Bei Yu</i>	
DREAM: A Data Replication Technique for Real-Time Mobile Ad-Hoc Network Databases	1206
<i>Prasanna Pabmanabhan, Le Gruenwald</i>	
An Air Index for Data Access over Multiple Wireless Broadcast Channels	1210
<i>Damdinsuren Amarmend, Masayoshi Aritsugi, Yoshinari Kanamori</i>	
Searching Local Information in Mobile Databases	1214
<i>Ouri Wolfson, Bo Xu, Huabei Yin, Hu Cao</i>	

POSTER SESSION 6: STREAM PROCESSING AND OPTIMIZATION

An Approach to Adaptive Memory Management in Data Stream Systems.....	1217
<i>Michael Cammert, Juergen Kraemer, Bernhard Seeger, Sonny Vaupel</i>	
Cluster Hull: A Technique for Summarizing Spatial Data Streams.....	1220
<i>John Hershberger, Nisheeth Shrivastava, Subhash Suri</i>	

Sliding Window Based Multi-way Join Algorithms over Distributed Data Streams	1223
<i>Dongdong Zhang, Jianzhong Li, Kimutai Kimeli, Weiping Wang</i>	
A Pipelined Framework for Online Cleaning of Sensor Data Streams	1226
<i>Shawn R. Jeffery, Gustavo Alonso, Michael J. Franklin, Wei Hong, Jennifer Widom</i>	
Revision Processing in a Stream Processing Engine: A High-Level Design.....	1229
<i>Esther Ryzkina, Anurag S. Maskey, Mitch Cherniack, Stan Zdonik</i>	
Approximating Streaming Window Joins Under CPU Limitations	1232
<i>Ahmed Ayad, Jeffrey Naughton, Stephen Wright, Utkarsh Srivastava</i>	

POSTER SESSION 7: SENSOR NETWORK AND SENSOR QUERIES

Monitoring Top-k Query in Wireless Sensor Networks	1235
<i>Minji Wu, Jianliang Xu, Xueyan Tang, Wang-Chien Lee</i>	
LB-Index: A Multi-Resolution Index Structure for Images	1238
<i>Vebjorn Ljosa, Arnab Bhattacharya, Ambuj K. Singh</i>	
Energy-Efficient Continuous Isoline Queries in Sensor Networks	1241
<i>Adam Silberstein, Rebecca Braynard, Jun Yang</i>	
Better Burst Detection	1245
<i>Xin Zhang, Dennis Shasha</i>	
Operators for Expensive Functions in Continuous Queries	1249
<i>Matthew Denny, Michael J. Franklin</i>	
Faster In-Network Evaluation of Spatial Aggregation in Sensor Networks	1252
<i>Dina Goldin</i>	
Threshold Similarity Queries in Large Time Series Databases	1256
<i>Johannes Aßfalg, Hans-Peter Kriegel, Peer Krüger, Peter Kunath, Alexey Pryakhin, Matthias Renz</i>	

POSTER SESSION 8: DATABASE SYSTEM INTERNALS, QUERY PROCESSING AND OPTIMIZATION

Automated Storage Management with QoS Guarantees	1260
<i>Lin Qiao, Balakrishna R. Iyer, Divyakant Agrawal, Amr El Abbadi</i>	
Materialized Sample Views for Database Approximation.....	1263
<i>Shantanu Joshi, Christopher Jermaine</i>	
Every Click You Make,I Will Be Fetching It: Efficient XML Query Processing in RDMS Using GUI-driven Prefetching	1266
<i>Sourav S Bhowmick, Sandeep Prakash</i>	
Achieving Class-Based QoS for Transactional Workloads	1269
<i>Bianca Schroeder, Mor Harchol-Balter, Arun Iyengar, Erich Nahum</i>	
Inferring a Serialization Order for Distributed Transactions	1272
<i>Khuzaima Daudjee, Kenneth Salem</i>	
Fractal Modeling of IP Network Traffic at Streaming Speeds.....	1275
<i>Flip Korn, S. Muthukrishnan, Yihua Wu</i>	
DEMO Chairs Message	1278
<i>Govi Govindarajan, Leo Mark, Arjen P. de Vries</i>	

DEMO SESSION 1: NEW APPROACHES IN DATA ENGINEERING

UNIDOOR: a Deductive Object-Oriented Database Management System	1280
<i>Mohammed K. Jaber, Andrei Voronkov</i>	
DaWaII: a Tool for the Integration of Autonomous Data Marts.....	1282
<i>Luca Cabibbo, Ivan Panella, Riccardo Torlone</i>	
ConQueSt: a Constraint-based Querying System for Exploratory Pattern Discovery	1283
<i>Francesco Bonchi, Fosca Giannotti, Claudio Lucchese, Salvatore, Orlando Raffaele Perego, Roberto Trasarti</i>	
MAPLE: A Mobile Scalable P2P Nearest Neighbor Query Model for Location-based Services.....	1285
<i>Wei-Shinn Ku, Roger Zimmerman, Chi-Ngai Wan, Haojun Wang</i>	
SIPPER: Selecting Informative Peers in Structured P2P Environment for Content-Based Retrieval.....	1286
<i>Shuigeng Zhou, Zheng Zhang, Aoying Zhou, Weining Qian</i>	
Simultaneous Pipelining in QPipe: Exploiting Work Sharing Opportunities Across Queries.....	1288
<i>Kun Gao, Stavros Harizopoulos, Ippokratis Pandis, Vladislav Shkapenyuk, Anastassia Ailamaki</i>	
Enabling Query Processing on Spatial Networks	1290
<i>Jagan Sankaranarayanan, Housman Alborzi, Hanan Samet</i>	
cgmOLAP: Efficient Parallel Generation and Querying of Terabyte Size ROLAP Data Cubes.....	1291
<i>Y. Chen, A. Rau-Chaplin, F. Dehne, T. Eavis, D. Green, E. Sithirasanen</i>	

DEMO SESSION 2: NEW DATA ENGINEERING APPLICATIONS

Practical Adaptation to Changing Resources in Grid Query Processing.....	1293
<i>Anastasios Gounaris, Norman W. Paton, Rizos Sakellariou, Alvaro A.A. Fernandes, Jim Smith, Paul Watson</i>	
ViEWNet: Visual Exploration of Region-Wide Traffic Networks.....	1295
<i>Hans-Peter Kriegel, Peter Kunath, Martin Pfeifle, Matthias Renz</i>	
AITVS: Advanced Interactive Traffic Visualization System	1297
<i>Chang-Tien Lu, Arnold P. Boedihardjo, Jinping Zheng</i>	
Stream Processing in Production-to-Business Software	1299
<i>Michael Cammert, Christoph Heinz, Jürgen Kraemer, Tobias Riemenschneider, Maxim Schwarzkopf, Bernhard Seeger, Alexander Zeiss</i>	
HSI: A Novel Framework for Efficient Automated Singer Identification in Large Music Databases	1301
<i>Jialie Shen, John Shepherd, Bin Cui, Kian-Lee Tan</i>	
XPlainer: An XPath Debugging Framework.....	1302
<i>Mariano P. Consens, John W.S. Liu, Bill O'Farrell</i>	
ACXESS - Access Control for XML with Enhanced Security Specifications.....	1303
<i>Sriram Mohan, Jonathan Klinginsmith, Arijit Sengupta, Yuqing Wu,</i>	
PETRANET: A Power Efficient Transaction Management Technique for Real-Time Mobile Ad-hoc Network Databases	1304
<i>Le Gruenwald, Percy Bernedo, Prasanna Padmanabhan</i>	

Author Index