

Table of Contents

Message From the Chairs.....	xvi
General Chairs.....	xix
Program Committee Members.....	xxi
External Reviewers.....	xxiv

RESEARCH SESSION 1: Views

MiniCount: Efficient Rewriting of COUNT-Queries Using Views.....	1
<i>Václav Lín, Vasilis Vassalos, and Prodromos Malakasiotis</i>	
Updates through Views: A New Hope.....	2
<i>Yannis Kotidis, Divesh Srivastava, and Yannis Velegarakis</i>	
Learning from Aggregate Views	3
<i>Bee-Chung Chen, Lei Chen, Raghu Ramakrishnan, and David R. Musicant</i>	

RESEARCH SESSION 2: Data Warehouse (1)

C-Cubing: Efficient Computation of Closed Cubes by Aggregation-Based Checking	4
<i>Dong Xin, Zheng Shao, Jiawei Han, and Hongyan Liu</i>	
A Primitive Operator for Similarity Joins in Data Cleaning.....	5
<i>Surajit Chaudhuri, Venkatesh Ganti, and Raghav Kaushik</i>	
Techniques for Warehousing of Sample Data.....	6
<i>Paul G. Brown and Peter J. Haas</i>	

RESEARCH SESSION 3: Query Processing and Uncertainty Reasoning

Working Models for Uncertain Data.....	7
<i>Anish Das Sarma, Omar Benjelloun, Alon Halevy, and Jennifer Widom</i>	
Reasoning about Approximate Match Query Results.....	8
<i>Sudipto Guha, Nick Koudas, Divesh Srivastava, Xiaohui Yu</i>	
The Gauss-Tree: Efficient Object Identification in Databases of Probabilistic Feature Vectors.....	9
<i>Christian Böhm, Alexey Pryakhin, Matthias Schubert</i>	

RESEARCH SESSION 4: Indexing and Optimization

Finding Fastest Paths on A Road Network with Speed Patterns.....	10
<i>Evangelos Kanoulas, Yang Du, Tian Xia, and Donghui Zhang</i>	
Approximation Techniques for Indexing the Earth Mover™'s Distance in Multimedia Databases.....	11
<i>Ira Assent, Andrea Wenning, and Thomas Seidl</i>	
Indexing for Dynamic Abstract Regions	12
<i>Joxan Jaffar, Roland H. C. Yap, and Kenny Q. Zhu</i>	

RESEARCH SESSION 5: XML and Semi-structured data

Efficient Processing of Updates in Dynamic XML Data	13
<i>Changqing Li, Tok Wang Ling, and Min Hu</i>	
A Complete and Efficient Algebraic Compiler for XQuery	14
<i>Christopher Ré, Jérôme Siméon, and Mary Fernández</i>	
Making Designer Schemas with Colors.....	15
<i>Nuwee Wiwatwattana, H.V. Jagadish, Laks V. S. Lakshmanan, and Divesh Srivastava</i>	

RESEARCH SESSION 6: Data Mining and Optimization

Mining Actionable Patterns by Role Models.....	16
<i>Ke Wang, Yuelong Jiang, and Alexander Tuzhilin</i>	
A Systematic Approach for Optimizing Complex Mining Tasks on Multiple Databases	17
<i>Ruoming Jin and Gagan Agrawal</i>	
New Sampling-Based Estimators for OLAP Queries	18
<i>Ruoming Jin, Leo Glimcher, Chris Jermaine, and Gagan Agrawal</i>	
Efficiently Evaluating Order Preserving Similarity Queries over Historical Market-Basket Data.....	19
<i>Reza Sherkat and Davood Rafiei</i>	

RESEARCH SESSION 7: Query Processing and Query Optimization

End-biased Samples for Join Cardinality Estimation	20
<i>Cristian Estan and Jeffrey F. Naughton</i>	
Laws for Rewriting Queries Containing Division Operators	21
<i>Ralf Rantzau and Christoph Mangold</i>	

R-trees with Update Memos.....	22
<i>Xiaopeng Xiong and Walid G. Aref</i>	
Compiled Query Execution Engine using JVM.....	23
<i>Jun Rao, Hamid Pirahesh, C. Mohan, and Guy Lohman</i>	

RESEARCH SESSION 8: Data Privacy and Security

l-Diversity: Privacy beyond K-Anonymity	24
<i>Ashwin Machanavajhala, Johannes Gehrke , Daniel Kifer, and Muthuramakrishnan Venkitasubramaniam</i>	
Mondrian Multidimensional K-Anonymity	25
<i>Kristen LeFevre, David J. DeWitt, and Raghu Ramakrishnan</i>	
Sovereign Joins.....	26
<i>Rakesh Agrawal, Dmitri Asonov, Murat Kantarcioglu, Yaping Li</i>	
Privacy Preserving Query Processing using Third Parties.....	27
<i>Fatih Emekci, Divyakant Agrawal, Amr El Abbadi, Aziz Gülboden</i>	

RESEARCH SESSION 9: Data Integration and Database Interoperability

Efficient Batch Top-k Search for Dictionary-Based Entity Recognition.....	28
<i>Amit Chandel, P. C. Nagesh, and Sunita Sarawagi</i>	
Integrating Unstructured Data into Relational Databases.....	29
<i>Imran R. Mansuri and Sunita Sarawagi</i>	
Clean Answers over Dirty Databases: A Probabilistic Approach.....	30
<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	31
<i>Ken Pu, Vagelis Hristidis, Nick Koudas</i>	
Hilda: A High-Level Language for Data-Driven Web Applications	32
<i>Fan Yang, Jayavel Shanmugasundaram, Mirek Riedewald, Johannes Gehrke, and Alan Demers</i>	

UNIT: User-centric Transaction Management in Web-Database Systems	33
<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.....	34
<i>H.V. Jagadish, Beng Chin Ooi, Quang Hieu Vu, Rong Zhang, Aoying Zhou</i>	
Transaction Time Support inside a Database Engine.....	35
<i>David Lomet, Roger Barga, Mohamed F. Mokbel, German Shegalov, Rui Wang, Yunyue Zhu</i>	
Expiration Times for Data Management	36
<i>Albrecht Schmidt, Christian S. Jensen, Simonas Saltenis</i>	

RESEARCH SESSION 12: Query Optimization and Data Structures

Scalable Exploration of Physical Database Design	37
<i>Arnd Christian König, Shubha U. Nabar</i>	
Closure-Tree: An Index Structure for Graph Queries.....	38
<i>Huahai He, Ambuj K. Singh</i>	
ISOMER: Consistent Histogram Construction Using Query Feedback.....	39
<i>Utkarsh Srivastava, Peter J. Haas, Volker Markl, Nimrod Megiddo, Marsel Kutsch, and 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	40
<i>Nikos Ntarmos, Peter Triantafillou, and Gerhard Weikum</i>	
Probabilistic Message Passing in Peer Data Management Systems	41
<i>Philippe Cudré-Mauroux, Karl Aberer, and Andras Feher</i>	
Approximating Aggregation Queries in Peer-to-Peer Networks	42
<i>Benjamin Arai, Gautam Das, Dimitrios Gunopulos, and Vana Kalogeraki</i>	
Distributed Evaluation of Continuous Equi-join Queries over Large Structured Overlay Networks	43
<i>Stratos Idreos, Christos Tryfonopoulos, and Manolis Koubarakis</i>	

RESEARCH SESSION 14: Web Queries

WebIQ: Learning from the Web to Match Deep-Web	44
<i>Wensheng Wu, AnHai Doan, Clement Yu</i>	
Answering Imprecise Queries over Autonomous Web Databases	45
<i>Ullas Nambiar and Subbarao Kambhampati</i>	
Merging Source Query Interfaces on Web Databases	46
<i>Eduard Dragut, Wensheng Wu, Prasad Sistla, Clement Yu, and Weiyi Meng</i>	
Query Selection Techniques for Efficient Crawling of Structured Web Sources	47
<i>Ping Wu, Ji-Rong Wen, Huan Liu, and Wei-Ying Ma</i>	

RESEARCH SESSION 15: Stream processing (1)

Approximate Data Collection in Sensor Networks using Probabilistic Models.....	48
<i>David Chu, Amol Deshpande, Joseph M. Hellerstein, and Wei Hong</i>	
Network-Aware Operator Placement for Stream-Processing Systems	49
<i>Peter Pietzuch, Jonathan Ledlie, Jeffrey Shneidman, Mema Roussopoulos, Matt Welsh, and Margo Seltzer</i>	
Unifying the Processing of XML Streams and Relational Data Streams	50
<i>Xin Zhou, Hetal Thakkar, and Carlo Zaniolo</i>	
Space-efficient Relative Error Order Sketch over Data Streams.....	51
<i>Ying Zhang, Xuemin Lin, Jian Xu, Flip Korn, and Wei Wang</i>	

RESEARCH SESSION 16: XML and XPath

Designing and Evaluating an XPath Dialect for Linguistic Queries	52
<i>Steven Bird, Yi Chen, Susan B. Davidson, Haejoong Lee, and Yifeng Zheng</i>	
Predicate-based Filtering of XPath Expressions	53
<i>Shuang Hou and H.-Arno Jacobsen</i>	
An Estimation System for XPath Expressions	54
<i>Hanyu Li, Mong Li Lee, Wynne Hsu, and Gao Cong</i>	

RESEARCH SESSION 17: Stream Processing (2)

Adaptive Reorganization of Coherency-Preserving Dissemination Tree for Streaming Data	55
<i>Yongluan Zhou, Beng Chin Ooi, Kian-Lee Tan, and Feng Yu</i>	

Declarative Network Monitoring with an Underprovisioned Query Processor.....	56
<i>Frederick Reiss and Joseph M. Hellerstein</i>	

What's Different: Distributed, Continuous Monitoring of Duplicate-Resilient Aggregates on Data Streams	57
<i>Graham Cormode, S. Muthukrishnan, and Wei Zhuang</i>	

RESEARCH SESSION 18: Database System Internals and Performance

Extending RDBMSs to Support Sparse Datasets Using an Interpreted Attribute Storage Format	58
<i>Jennifer L. Beckmann, Alan Halverson, Rajasekar Krishnamurthy, and Jeffrey F. Naughton</i>	

Super-Scalar RAM-CPU Cache Compression	59
<i>Marcin Zukowski, Sándor Héman, Niels Nes, and Peter Boncz</i>	

How to Determine a Good Multi-Programming Level for External Scheduling	60
<i>Bianca Schroeder, Mor Harchol-Balter, Arun Iyengar, Erich Nahum, and Adam Wierman</i>	

RESEARCH SESSION 19: XML Databases and XML Query Optimization

XSeed: Accurate and Fast Cardinality Estimation for XPath Queries	61
<i>Ning Zhang, Tamer Ozsu, Ashraf Aboulmaga, and Ihab F. Ilyas</i>	

Estimating XML Structural Join Size Quickly and Economically	62
<i>Cheng Luo, Zhewei Jiang, Wen-Chi Hou, Feng Yan, and Chih-Fang Wang</i>	

XCluster Synopses for Structured XML Content	63
<i>Neoklis Polyzotis, and Minos Garofalakis</i>	

RESEARCH SESSION 20: Skyline Query Processing

Robust Cardinality and Cost Estimation for Skyline Operator	64
<i>Surajit Chaudhuri, Nilesh Dalvi, and Raghav Kaushik</i>	

SUBSKY: Efficient Computation of Skylines in Subspaces	65
<i>Yufei Tao, Xiaokui Xiao, and Jian Pei</i>	

Skyline queries against mobile lightweight devices in MANETs.....	66
<i>Zhiyong Huang, Christian S. Jensen, Hua Lu, and Beng Chin Ooi</i>	

RESEARCH SESSION 21 Query Management

Approximately Processing Multi-granularity Aggregate Queries over a Data Stream	67
Shouke Qin, Weining Qian, and Aoying Zhou	
A Sampling-Based Approach to Optimizing Top-k Queries in Sensor Networks.....	68
<i>Adam Silberstein, Rebecca Braynard, Carla Ellis, Kamesh Munagala, and Jun Yang</i>	
Précis: The Essence of a Query Answer	69
<i>Georgia Koutrika, Alkis Simitsis, and Yannis Ioannidis</i>	

RESEARCH SESSION 22: Ubiquitous Data Management

Processing Window Queries in Wireless Sensor Networks.....	70
<i>Yingqi Xu, Wang-Chien Lee, Jianliang Xu, and Gail Mitchell</i>	
Effective Density Queries of Continuously Moving Objects.....	71
<i>Christian S. Jensen, Dan Lin, Beng Chin Ooi, and Rui Zhang</i>	
Efficient Aggregation of Ranked Inputs.....	72
<i>Nikos Mamoulis, Kit Hung Cheng, Man Lung Yiu, and David W. Cheung</i>	

RESEARCH SESSION 23: Graph Databases and Algorithms

CLAN: An Algorithm for Mining Closed Cliques from Large Dense Graph Databases	73
<i>Jianyong Wang, Zhiping Zeng, and Lizhu Zhou</i>	
A Partition-Based Approach to Graph Mining	74
<i>Junmei Wang, Wynne Hsu, Mong Li Lee, and Chang Sheng</i>	
Dual Labeling: Answering Graph Reachability Queries in Constant Time.....	75
<i>Haixun Wang, Hao He, Jun Yang, Philip S. Yu, and Jeffrey Xu Yu</i>	

RESEARCH SESSION 24: Nearest Neighbor Search in Temporal and Spatial Databases

Reverse Nearest Neighbors Search in Ad-hoc Subspaces.....	76
<i>Man Lung Yiu and Nikos Mamoulis</i>	
Continuous Reverse Nearest Neighbor Monitoring	77
<i>Tian Xia and Donghui Zhang</i>	
Surface KNN Query Processing.....	78
<i>Ke Deng, Xiaofang Zhou, Heng Tao Shen, Kai Xu, and Xuemin Lin</i>	

RESEARCH SESSION 25: Stream Query Optimization

An Efficient XPath Query Processor for XML Streams	79
<i>Yi Chen, Susan B. Davidson, and Yifeng Zheng</i>	
SketchTree: Approximate Tree Pattern Counts over Streaming Labeled Trees	80
<i>Praveen Rao and Bongki Moon</i>	
Characterizing and Exploiting Reference Locality in Data Stream Applications	81
<i>Feifei Li, Ching Chang, George Kollios, and Azer Bestavros</i>	

RESEARCH SESSION 26: Advanced Query Processing

MONDRIAN: Annotating and querying databases through colors and blocks	82
<i>Floris Geerts, Anastasios Kementsietsidis, and Diego Milano</i>	
Warehousing and Analyzing Massive RFID Data Sets	83
<i>Hector Gonzalez, Jiawei Han, Xiaolei Li, and Diego Klabjan</i>	

RESEARCH SESSION 27: Temporal and Spatial Data Management (2)

Query Recomposition: A Multiple Neighborhood Approach to Relevance Feedback Processing in Content-Based Image Retrieval.....	84
<i>Kien A. Hua, Ning Yu, and Danzhou Liu</i>	
Nearest Surrounding Queries	85
<i>Ken C. K. Lee, Wang-Chien Lee, and Hong Va Leong</i>	
Closest-Point-of-Approach Join for Moving Object Histories	86
<i>Subramanian Arumugam and Christopher Jermaine</i>	

RESEARCH SESSION 28: Scientific and Biological Databases and Bioinformatics

Declarative Querying for Biological Sequences	87
<i>Sandeep Tata, Jignesh M. Patel, James S. Friedman, and Anand Swaroop</i>	
Searching Substructures with Superimposed Distance.....	88
<i>Xifeng Yan, Feida Zhu, Jiawei Han, and Philip S. Yu</i>	
Mining Shifting-and-Scaling Co-Regulation Patterns on Gene Expression Profiles	89
<i>Xin Xu, Ying Lu, Anthony K.H. Tung, and Wei Wang</i>	

INDUSTRY SESSION 1: Databases, Services, and Processes

AutoGlobe: An Automatic Administration Concept for Service-Oriented Database Applications	90
<i>Stefan Seltzscham, Daniel Gmach, Stefan Krompass, and Alfons Kemper</i>	

Collaborative Business Process Support in IHE XDS through ebXML Business Process.....	91
<i>Asuman Dogac, Veli Bicer, and Alper Okcan</i>	

Taming Compliance with Sarbanes-Oxley Internal Controls Using Database Technology.....	92
<i>Rakesh Agrawal, Christopher Johnson, Jerry Kiernan, and Frank Leymann</i>	

INDUSTRY SESSION 2: RDF, Ontologies, Metadata

RDF Object Type and Reification in the Database	93
<i>Nicole Alexander and Siva Ravada</i>	

RDF/RDFS-based Relational Database Integration.....	94
<i>Huajun Chen, Zhaohui Wu, Heng Wang, and Yuxin Mao</i>	

Supporting Keyword Columns with Ontology-based Referential Constraints in DBMS.....	95
<i>Eugene Chong, Souripriya Das, George Eadon, and Jagannathan Srinivasan</i>	

Experiment Management with Metadata-based Integration for Collaborative Scientific Research	96
<i>Fusheng Wang, Peiya Liu, John Pearson, Fred Azar, and Gerald Madlmayr</i>	

INDUSTRY SESSION 3: Querying and Logging

SQL to XQuery Translation in the AquaLogic Data Services Platform	97
<i>Sunil Jigyasu, Sujeet Banerjee, Vinayak Borkar, Michael Carey, Kanad Dixit, Anil Malkani, and Sachin Thatte</i>	

On Pushing Multilingual Query Operators into Relational Engines.....	98
<i>A Kumaran, Pavan K. Chowdary, and Jayant R. Haritsa</i>	

SIREN: A Memory-Conserving, Snapshot-Consistent Checkpoint Algorithm for in-Memory Databases	99
<i>Antti-Pekka Lienes and Antoni Wolski</i>	

INDUSTRY SESSION 4: Data Warehousing and Analysis

Space-Partitioning Trees in PostgreSQL: Realization and Performance	100
<i>M.Y. Eitabakh, Ramy Eltarras, and Walid G. Aref</i>	

Automatic Sales Lead Generation from Web Data	101
<i>Ganesh Ramakrishnan, Sachindra Joshi, Sumit Negi, Raghu Krishnapuram, and Sreeram Balakrishnan</i>	

Load Balancing for Multi-tiered Database Systems through Autonomic Placement of Materialized Views	102
<i>Wen-Syan Li, Daniel C. Zilio, Vishal Batra, Mahadevan Subramanian, Calisto Zuzarte, and Inderpal Narang</i>	

ADVANCED SEMINAR 1: Schema and Data Translation	103
<i>Paolo Atzeni</i>	

ADVANCED SEMINAR 2: Foundations of Automated Database Tuning	104
<i>Surajit Chaudhuri, Gerhard Weikum</i>	

ADVANCED SEMINAR 3: Models and Methods for Privacy-Preserving Data Publishing and Analysis	105
<i>Johannes Gehrke</i>	

ADVANCED SEMINAR 4 Mining, Indexing, and Similarity Search in Graphs and Complex Structures	106
<i>Jiawei Han, Xifeng Yan, and Philip S. Yu</i>	

ADVANCED SEMINAR 5 Query co-Processing on Commodity Hardware	107
<i>Anastassia Ailamaki, Naga K. Govindaraju, and Dinesh Manocha</i>	

POSTER SESSION 1: Data Integration, Data Mining, Data Warehousing

Efficient Continuous Skyline Computation	108
<i>Michael Morse, Jignesh M. Patel, and William I. Grosky</i>	
Detecting Duplicates in Complex XML Data	109
<i>Melanie Weis and Felix Naumann</i>	
SaveRF: Towards Efficient Relevance Feedback Search	110
<i>Heng Tao Shen, Beng Chin Ooi, and Kian-Lee Tan</i>	
On the Inverse Classification Problem and its Applications	111
<i>Charu C. Aggarwal, Chen Chen, and Jiawei Han</i>	
MIC Framework: An Information-Theoretic Approach to Quantitative Association Rule Mining	112
<i>Yiping Ke, James Cheng, and Wilfred Ng</i>	

Efficient Discovery of Emerging Frequent Patterns in Arbitrary Windows on Data Streams	113
<i>Xiaoming Jin, Xinqiang Zuo, Kwok-Yan Lam, Jianmin Wang, and Jiaguang Sun</i>	

Top-down Mining of Interesting Patterns from Very High Dimensional Data	114
<i>Hongyan Liu, Jiawei Han, Dong Xin, and Zheng Shao</i>	

Mining Dense Periodic Patterns in Time Series Data.....	115
<i>Chang Sheng, Wynne Hsu, and Mong Li Lee</i>	

POSTER SESSION 2: Data Privacy and Security

Private Updates to Anonymous Databases	116
<i>Alberto Trombetta and Elisa Bertino</i>	

Provable Security for Outsourcing Database Operations	117
<i>Sergei Evdokimov, Matthias Fischmann, and Oliver Günther</i>	

Composition and Disclosure of Unlinkable Distributed Databases	118
<i>Bradley Malin and Latanya Sweeney</i>	

A Technique for Optimal Adaptation of Time-Dependant Workflow with Security.....	119
<i>Basit Shafiq, Arjmand Samuel, Elisa Bertino, and Arif Ghafoor</i>	

POSTER SESSION 3: Web Data Management and Workflow

Segmentation of Publication Records of Authors from the Web	120
<i>Wei Zhang, Neil Smalheiser, Vette Torvik, and Clement Yu</i>	

Text Classification Improved through Automatically Extracted Sequences	121
<i>Dou Shen, Jian-Tao Sun, Qiang Yang, Hui Zhao, and Zheng Chen</i>	

Holistic Query Interface Matching using Parallel Schema Matching	122
<i>Weifeng Su, Jiying Wang, and Frederick Lochovsky</i>	

Extracting Objects from the Web.....	123
<i>Zaiqing Nie, Fei Wu, Ji-Rong Wen, and Wei-Ying Ma</i>	

On the Controlled Evolution of Process Choreographies	124
<i>Stefanie Rinderle, Andreas Wombacher, and Manfred Reichert</i>	

Automating the Design and Construction of Query Forms	125
<i>Magesh Jayapandian and H.V. Jagadish</i>	

POSTER SESSION 4: XML Data Management

U-Filter: A Lightweight XML View Update Checker	126
<i>Ling Wang, Elke A. Rudensteiner, and Murali Mani</i>	

XCut: Indexing XML Data for Efficient Twig Evaluation	127
<i>Simon Sheu, and Nigel Wu</i>	

Algebraic Optimization of Nested XPath Expressions	128
<i>Matthias Brantner, Sven Helmer, Carl-Christian Kanne, and Guido Moerkotte</i>	
Incremental Maintenance of Materialized XQuery Views	129
<i>Maged El-Sayed, Elke A. Rudensteiner, and Murali Mani</i>	
Don't be a Pessimist: Use Snapshot based Concurrency Control for XML.....	130
<i>Zeeshan Sardar and Bettina Kemme</i>	
Using XML to Build Efficient Transaction-Time Temporal Database Systems on Relational Databases	131
<i>Fusheng Wang, Xin Zhou, and Carlo Zaniolo</i>	

POSTER SESSION 5: Peer to Peer and Mobile Data Management

Partial Selection Query in Peer-to-Peer Databases	132
<i>Farnoush Banaei-Kashani, and Cyrus Shahabi</i>	
HiWaRPP --- Hierarchical Wavelet-based Retrieval on Peer-to-Peer Network	133
<i>Mihai Lupu, Bei Yu</i>	
DREAM: A Data Replication Technique for Real-Time Mobile Ad-Hoc Network Databases.....	134
<i>Prasanna Pabmanabhan, Le Gruenwald</i>	
An Air Index for Data Access over Multiple Wireless Broadcast Channels	135
<i>Damdinsuren Amarmend, Masayoshi Aritsugi, Yoshinari Kanamori</i>	
Searching Local Information in Mobile Databases.....	136
<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	137
<i>Michael Cammert, Juergen Kraemer, Bernhard Seeger, and Sonny Vaupel</i>	
Cluster Hull: A Technique for Summarizing Spatial Data Streams	138
<i>John Hershberger, Nisheeth Shrivastava, and Subhash Suri</i>	
Sliding Window Based Multi-way Join Algorithms over Distributed Data Streams	139
<i>Dongdong Zhang, Jianzhong Li, Kimutai Kimeli, and Weiping Wang</i>	
A Pipelined Framework for Online Cleaning of Sensor Data Streams.....	140
<i>Shawn R. Jeffery, Gustavo Alonso, Michael J. Franklin, Wei Hong, and Jennifer Widom</i>	
Revision Processing in a Stream Processing Engine: A High-Level Design	141
<i>Esther Ryvkina, Anurag S. Maskey, Mitch Cherniack, and Stan Zdonik</i>	

Approximating Streaming Window Joins Under CPU Limitations	142
<i>Ahmed Ayad, Jeffrey Naughton, Stephen Wright, and Utkarsh Srivastava</i>	

POSTER SESSION 7: Sensor Network and Sensor Queries

Monitoring Top-k Query in Wireless Sensor Networks	143
<i>Minji Wu, Jianliang Xu, Xueyan Tang, and Wang-Chien Lee</i>	
LB-Index: A Multi-Resolution Index Structure for Images.....	144
<i>Vebjorn Ljosa, Arnab Bhattacharya, and Ambuj K. Singh</i>	
Energy-Efficient Continuous Isoline Queries in Sensor Networks	145
<i>Adam Silberstein, Rebecca Braynard, and Jun Yang</i>	
Better Burst Detection.....	146
<i>Xin Zhang and Dennis Shasha</i>	
Operators for Expensive Functions in Continuous Queries	147
<i>Matthew Denny and Michael J. Franklin</i>	
Faster In-Network Evaluation of Spatial Aggregation in Sensor Networks.....	148
<i>Dina Goldin</i>	
Threshold Similarity Queries in Large Time Series Databases	149
<i>Johannes Aßfalg, Hans-Peter Kriegel, Peer Krüger, Peter Kunath, Alexey Pryakhin, and Matthias Renz</i>	

POSTER SESSION 8: Database System Internals, Query Processing and Optimization

Automated Storage Management with QoS Guarantees	150
<i>Lin Qiao, Balakrishna R. Iyer, Divyakant Agrawal, and Amr El Abbadi</i>	
Materialized Sample Views for Database Approximation	151
<i>Shantanu Joshi and Christopher Jermaine</i>	
Every Click You Make,I Will Be Fetching It: Efficient XML Query Processing in RDMS Using GUI-driven Prefetching.....	152
<i>Sourav S Bhowmick and Sandeep Prakash</i>	
Achieving Class-Based QoS for Transactional Workloads	153
<i>Bianca Schroeder, Mor Harchol-Balter, Arun Iyengar, and Erich Nahum</i>	
Inferring a Serialization Order for Distributed Transactions.....	154
<i>Khuzaima Daudjee and Kenneth Salem</i>	

Fractal Modeling of IP Network Traffic at Streaming Speeds.....	155
<i>Flip Korn, S. Muthukrishnan, and Yihua Wu</i>	

DEMO CHAIRS MESSAGE.....	156
<i>Govi Govindarajan, Leo Mark, and Arjen P. de Vries</i>	

DEMO SESSION 1: New Approaches in Data Engineering

UNIDOOR: a Deductive Object-Oriented Database Management System.....	157
<i>Mohammed K. Jaber and Andrei Voronkov</i>	

DaWaII: a Tool for the Integration of Autonomous Data Marts	158
<i>Luca Cabibbo, Ivan Panella, and Riccardo Torlone</i>	

ConQueSt: a Constraint-based Querying System for Exploratory Pattern Discovery	159
<i>Francesco Bonchi, Fosca Giannotti, Claudio Lucchese, Salvatore, Orlando Raffaele Perego, and Roberto Trasarti</i>	

MAPLE: A Mobile Scalable P2P Nearest Neighbor Query Model for Location-based Services	160
<i>Wei-Shinn Ku, Roger Zimmerman, Chi-Ngai Wan, and Haojun Wang</i>	

SIPPER: Selecting Informative Peers in Structured P2P Environment for Content-Based Retrieval	161
<i>Shuigeng Zhou, Zheng Zhang, Aoying Zhou, and Weining Qian</i>	

Simultaneous Pipelining in QPipe: Exploiting Work Sharing Opportunities Across Queries	162
<i>Kun Gao, Stavros Harizopoulos, Ippokratis Pandis, Vladislav Shkapenyuk, and Anastassia Ailamaki</i>	

Enabling Query Processing on Spatial Networks	163
<i>Jagan Sankaranarayanan, Houman Alborzi, and Hanan Samet</i>	

cgmOLAP: Efficient Parallel Generation and Querying of Terabyte Size ROLAP Data Cubes	164
<i>Y. Chen, A. Rau-Chaplin, F. Dehne, T. Eavis, D. Green, E. Sithirasenan</i>	

DEMO SESSION 2: New Data Engineering Applications

Practical Adaptation to Changing Resources in Grid Query Processing	165
<i>Anastasios Gounaris, Norman W. Paton, Rizos Sakellariou, Alvaro A.A. Fernandes, Jim Smith, and Paul Watson</i>	

ViEWNet: Visual Exploration of Region-Wide Traffic Networks.....	166
<i>Hans-Peter Kriegel, Peter Kunath, Martin Pfeifle, and Matthias Renz</i>	

AITVS: Advanced Interactive Traffic Visualization System	167
<i>Chang-Tien Lu, Arnold P. Boedihardjo, and Jinping Zheng</i>	

Stream Processing in Production-to-Business Software	168
<i>Michael Cammert, Christoph Heinz, Jüergen Kraemer, Tobias Riemenschneider,</i> <i>Maxim Schwarzkopf, Bernhard Seeger, and Alexander Zeiss</i>	
HSI: A Novel Framework for Efficient Automated Singer Identification in Large Music Databases.....	169
<i>Jialie Shen, John Shepherd, Bin Cui, and Kian-Lee Tan</i>	
XPlainer: An XPath Debugging Framework.....	170
<i>Mariano P. Consens, John W.S. Liu, and Bill O'Farrell</i>	
ACXESS - Access Control for XML with Enhanced Security Specifcations	171
<i>Sriram Mohan,Jonathan Klinginsmith, Arijit Sengupta, and Yuqing Wu,</i>	
PETRANET: A Power Efficient Transaction Management Technique for Real-Time Mobile Ad-hoc Network Databases.....	172
<i>Le Gruenwald, Percy Bernedo, and Prasanna Padmanabhan,</i>	