

Table of Contents

ii Copyright Notice	xi Program Committees
iii Welcome from the General Chairs	xxi External Reviewers
v Message from the Program Chairs	xxviii VLDB Endowment
vii Organizing Committee	xxix Table of Contents

Keynotes

1

Enabling Real Time Data Analysis.

Divesh Srivastava (AT&T Labs Research, United States of America), Lukasz Golab (AT&T Labs Research, United States of America), Rick Greer (AT&T Labs Research, United States of America), Theodore Johnson (AT&T Labs Research, United States of America), Joseph Seidel (AT&T Labs Research, United States of America), Vladislav Shkapenyuk (AT&T Labs Research, United States of America), Oliver Spatscheck (AT&T Labs Research, United States of America), Jennifer Yates (AT&T Labs Research, United States of America).

3

High-End Biological Imaging Generates Very Large 3D+ and Dynamic Datasets

Paul Matsudaira (National University of Singapore, Republic of Singapore).

10-Year Best Paper Awards

Keynote Sessions

4

Dealing with Web Data: History and Look ahead

Junghoo Cho (University of California Los Angeles, United States of America), Hector Garcia-Molina (Stanford University, United States of America).

5

Database Replication: a Tale of Research across Communities

Bettina Kemme (McGill University, Canada), Gustavo Alonso (Eidgenössische Technische Hochschule Zürich, Switzerland).

Research Sessions

Database Security

13

Building Disclosure Risk Aware Query Optimizers for Relational Databases

Mustafa Canim (University of Texas at Dallas, United States of America), Murat Kantarcioglu (University of Texas at Dallas, United States of America), Bijit Hore (University of California Irvine, United States of America), Sharad Mehrotra (University of California Irvine, United States of America).

25

Secure Personal Data Servers: a Vision Paper

Tristan Allard (University of Versailles, France), Nicolas Anceaux (Institut National de Recherche en Informatique et Automatique, France), Luc Bouganim (Institut National de Recherche en Informatique et Automatique, France), Yanli Guo (Institut National de Recherche en Informatique et Automatique, France), Lionel Le Folgoc (Institut National de Recherche en Informatique et Automatique, France), Benjamin Nguyen (Institut National de Recherche en Informatique et Automatique, France), Philippe Pucheral (Institut National

de Recherche en Informatique et Automatique, France), Indrajit Ray (Colorado State University, United States of America), Indrakshi Ray (Colorado State University, United States of America), Shaoyi Yin (Institut National de Recherche en Informatique et Automatique, France).

36

PolicyReplay: Misconfiguration-Response Queries for Data Breach Reporting

Daniel Fabbri (University of Michigan, United States of America), Kristen LeFevre (University of Michigan, United States of America), Qiang Zhu (University of Michigan, United States of America).

Parallel and Distributed Databases

48

Schism: a Workload-Driven Approach to Database Replication and Partitioning

Carlo Curino (Massachusetts Institute of Technology, United States of America), Yang Zhang (Massachusetts Institute of Technology, United States of America), Evan Jones (Massachusetts Institute of Technology, United States of America), Samuel Madden (Massachusetts Institute of Technology, United States of America).

58

Ten Thousand SQLs: Parallel Keyword Queries Computing

Lu Qin (The Chinese University of Hong Kong, People's Republic of China), Jefferey Yu (The Chinese University of Hong Kong, People's Republic of China), Lijun Chang (The Chinese University of Hong Kong, People's Republic of China).

70

The Case for Determinism in Database Systems

Alexander Thomson (Yale University, United States of America), Daniel Abadi (Yale University, United States of America).

Data Exchange

81

MapMerge: Correlating Independent Schema Mappings

Bogdan Alexe (University of California Santa Cruz, United States of America), Mauricio Hernández (IBM Research, United States of America), Lucian Popa (IBM Almaden Research Center, United States of America), Wang-Chiew Tan (University of California Santa Cruz, United States of America).

93

Chase Termination: A Constraints Rewriting Approach

Francesca Spezzano (Università della Calabria, Italy), Sergio Greco (Università della Calabria, Italy).

105

Scalable Data Exchange with Functional Dependencies

Bruno Marnette (University of Oxford, United Kingdom), Giansalvatore Mecca (Università della Basilicata, Italy), Paolo Papotti (Università Roma Tre, Italy).

Database Services and Applications

117

Interactive Route Search in the Presence of Order Constraints

Roy Levin (Technion-Israel Institute of Technology, Israel), Yaron Kanza (Technion-Israel Institute of Technology, Israel), Eliyahu Safra, Yehoshua Sagiv (Hebrew University of Jerusalem, Israel).

129

Energy Management for MapReduce Clusters

Willis Lang (University of Wisconsin-Madison,

United States of America), Jignesh Patel (University of Wisconsin-Madison, United States of America).

140

Toward Scalable Keyword Search over Relational Data

Akanksha Baid (University of Wisconsin-Madison, United States of America), Ian Rae (University of Wisconsin-Madison, United States of America), Jiexing Li (University of Wisconsin-Madison, United States of America), AnHai Doan (University of Wisconsin-Madison, United States of America), Jeffrey Naughton (University of Wisconsin-Madison, United States of America).

Data Models and Languages

150

From Regular Expressions to Nested Words: Unifying Languages and Query Execution for Relational and XML Sequences

Barzan Mozafari (University of California Los Angeles, United States of America), Kai Zeng (University of California Los Angeles, United States of America), Carlo Zaniolo (University of California Los Angeles, United States of America).

162

Avalanche-Safe LINQ Compilation

Torsten Grust (Universität Tübingen, Germany), Jan Rittinger (Universität Tübingen, Germany), Tom Schreiber (Universität Tübingen, Germany).

173

Towards Certain Fixes with Editing Rules and Master Data

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (Harbin Institute of Technology, People's Republic of China), Shuai Ma (University of Edinburgh, United Kingdom), Nan Tang (University of Edinburgh,

United Kingdom), Wenyan Yu (University of Edinburgh, United Kingdom).

Semantics

185

Explaining Missing Answers to SPJUA Queries

Melanie Herschel (Universität Tübingen, Germany), Mauricio Hernández (IBM Research, United States of America).

197

Sampling the Repairs of Functional Dependency Violations under Hard Constraints

George Beskales (University of Waterloo, Canada), Ihab Ilyas (University of Waterloo, Canada), Lukasz Golab (AT&T Labs - Research, United States of America).

208

Evaluating Entity Resolution Results

David Menestrina (Stanford University, United States of America), Steven Whang (Stanford University, United States of America), Hector Garcia-Molina (Stanford University, United States of America).

Stream Databases

220

High-Performance Dynamic Pattern Matching over Disordered Streams

Badrish Chandramouli (Microsoft Research, United States of America), Jonathan Goldstein (Microsoft Research, United States of America), David Maier (Portland State University, United States of America).

232

SECRET: A Model for Analysis of the Execution Semantics of Stream Processing Systems

Irina Botan (Eidgenössische Technische Hochschule Zürich, Switzerland), Roozbeh Derakhshan (Eidgenössische Technische Hochschule

Zürich, Switzerland), Nihal Dindar (Eidgenössische Technische Hochschule Zürich, Switzerland), Laura Haas (IBM Almaden Research Center, United States of America), Renée Miller (University of Toronto, Canada), Nesime Tatbul (Eidgenössische Technische Hochschule Zürich, Switzerland).

244

Recognizing Patterns in Streams with Imprecise Timestamps

Haopeng Zhang (University of Massachusetts Amherst, United States of America), Yanlei Diao (University of Massachusetts Amherst, United States of America), Neil Immerman (University of Massachusetts Amherst, United States of America).

RDF and Graphs

256

x-RDF-3X: Fast Querying, High Update Rates, and Consistency for RDF Databases

Thomas Neuman (Technische Universität München, Germany), Gerhard Weikum (Max-Planck Institute for Informatics, Germany).

264

Graph Pattern Matching: From Intractable to Polynomial Time

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (Harbin Institute of Technology, People's Republic of China), Shuai Ma (University of Edinburgh, United Kingdom), Nan Tang (University of Edinburgh, United Kingdom), Yinghui Wu (University of Edinburgh, United Kingdom), Yunpeng Wu (National University of Defense Technology, People's Republic of China).

276

GRAIL: Scalable Reachability Index for Large Graphs

Hilmi Yildirim (Rensselaer Polytechnic Institute, United States of America), Vineet Chaoji

(Yahoo! Research Labs, United States of America), Mohammed Zaki (Rensselaer Polytechnic Institute, United States of America).

Middleware Platforms for Data Management

285

HaLoop: Efficient Iterative Data Processing on Large Clusters

Yingyi Bu (University of Washington, United States of America), Bill Howe (University of Washington, United States of America), Magdalena Balazinska (University of Washington, United States of America), Michael Ernst (University of Washington, United States of America).

297

The Impact of Virtual Views on Containment

Michael Benedikt (University of Oxford, United Kingdom), Georg Gottlob (University of Oxford, United Kingdom).

309

Updatable and Evolvable Transforms for Virtual Databases

James Terwilliger (Microsoft Corporation, United States of America), Lois Delcambre (Portland State University, United States of America), David Maier (Portland State University, United States of America), Jeremy Steinhauer (Portland State University, United States of America), Scott Britell (Portland State University, United States of America).

Novel/Advanced Applications

320

Navigating in Complex Mashed-Up Applications

Daniel Deutch (Tel Aviv University, Israel), Ohad Greenshpan (Tel Aviv University, Israel), Tova Milo (Tel Aviv University, Israel).

330

Dremel: Interactive Analysis of Web-Scale Datasets

Sergey Melnik (Google, United States of America), Andrey Gubarev (Google, United States of America), Jing Jing Long (Google, United States of America), Geoffrey Romer (Google, United States of America), Shiva Shivakumar (Google, United States of America), Matt Tolton (Google, United States of America), Theo Vassilakis (Google, United States of America).

340

On Graph Query Optimization in Large Networks

Peixiang Zhao (University of Illinois at Urbana-Champaign, United States of America), Jiawei Han (University of Illinois at Urbana-Champaign, United States of America).

Ranking Queries

352

Proximity Rank Join

Davide Martinenghi (Politecnico di Milano, Italy), Marco Tagliasacchi (Politecnico di Milano, Italy).

364

Identifying the Most Influential Data Objects with Reverse Top-k Queries

Akrivi Vlachou (Norwegian University of Science and Technology, Norway), Christos Doukolidis (Norwegian University of Science and Technology, Norway), Kjetil Nørvag (Norwegian University of Science and Technology, Norway), Yannis Kotidis (Athens University of Economics and Business, Greece).

373

Retrieving Top-k Prestige-Based Relevant Spatial Web Objects

Xin Cao (Nanyang Technological University, Republic of Singapore), Gao Cong (Nanyang Technological University, Republic of Singapore), Christian Jensen (Aarhus University, Denmark).

Spatial and Temporal Databases

385

Parsimonious Linear Fingerprinting for Time Series

Lei Li (Carnegie Mellon University, United States of America), B. Aditya Prakash (Carnegie Mellon University, United States of America), Christos Faloutsos (Carnegie Mellon University, United States of America).

397

The HV-tree: a Memory Hierarchy Aware Version Index

Rui Zhang (University of Melbourne, Australia), Martin Stradling (University of Melbourne, Australia).

409

Transforming Range Queries To Equivalent Box Queries To Optimize Page Access

Sakti Pramanik (Michigan State University, United States of America), Alok Watve (Michigan State University, United States of America), Chad Meiners (Michigan State University, United States of America), Alex Liu (Michigan State University, United States of America).

Record Linkage

417

Record Linkage with Uniqueness Constraints and Erroneous Values

Songtao Guo (AT&T Interactive, United States of America), Xin Dong (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America), Remi Zajac (AT&T Interactive, United States of America).

429

On-the-Fly Entity-Aware Query Processing in the Presence of Linkage

Ekaterini Ioannou (L3S Research Center, Germany), Wolfgang Nejdl (L3S Research Center, Germany), Claudia Niederée (L3S Research Center, Germany), Yannis Velegrakis (University of Trento, Italy).

439

Behavior Based Record Linkage

Mohamed Yakout (Purdue University, United States of America), Ahmed Elmagarmid (Purdue University, United States of America), Hazem Elmeleegy (Purdue University, United States of America), Mourad Ouzzani (Purdue University, United States of America), Alan Qi (Purdue University, United States of America).

Experimental Analysis and Performance

449

iGraph: A Framework for Comparisons of Disk-Based Graph Indexing Techniques

Wook-Shin Han (Kyungpook National University, Republic of Korea), Jinsoo Lee (Kyungpook National University, Republic of Korea), Minh-Duc Pham (Kyungpook National University, Republic of Korea), Jeffrey Yu (The Chinese University of Hong Kong, People's Republic of China).

460

Runtime Measurements in the Cloud: Observing, Analyzing, and Reducing Variance

Jörg Schad (Saarland University, Germany), Jens Dittrich (Saarland University, Germany), Jorge-Arnulfo Quiané-Ruiz (Saarland University, Germany).

472

The Performance of MapReduce: An In-depth Study

Dawei Jiang (National University of Singapore, Republic of Singapore), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Lei Shi (National University of Singapore, Republic of Singapore), Sai Wu (National University of Singapore, Republic of Singapore).

484

Evaluation of entity resolution approaches on real-world match problems

Hanna Köpcke (University of Leipzig, Germany), Andreas Thor (University of Leipzig, Germany), Erhard Rahm (University of Leipzig, Germany).

Cloud Computing

494

MRShare: Sharing Across Multiple Queries in MapReduce

Tomasz Nykiel (University of Toronto, Canada), Michalis Potamias (Boston University, United States of America), Chaitanya Mishra (Facebook, United States of America), George Kollios (Boston University, United States of America), Nick Koudas (University of Toronto, Canada).

506

Towards Elastic Transactional Cloud Storage with Range Query Support

Hoang Tam Vo (National University of Singapore, Republic of Singapore), Chun Chen (Zhejiang University, People's Republic of China), Beng Chin Ooi (National University of Singapore, Republic of Singapore).

518

Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)

Jens Dittrich (Saarland University, Germany), Jorge-Arnulfo Quiané-Ruiz (Saarland University, Germany), Alekh Jindal (Saarland University/ International Max Planck Research

School for Computer Science, Germany), Yagiz Kargin (International Max Planck Research School for Computer Science, Germany), Vinay Setty (International Max-Planck Institute for Informatics, Germany), Jörg Schad (Saarland University, Germany).

Query Processing I

530

Slicing Long-Running Queries

Nicolas Bruno (Microsoft Research, United States of America), Vivek Narasayya (Microsoft Research, United States of America), Ravi Ramamurthy (Microsoft Research, United States of America).

542

Sharing-Aware Horizontal Partitioning for Exploiting Correlations During Query Processing

Kostas Tzoumas (Aalborg University, Denmark), Amol Deshpande (University of Maryland College Park, United States of America), Christian Jensen (Aalborg University, Denmark).

554

Advanced Processing for Ontological Queries

Andrea Cali (University of Oxford, United Kingdom), Georg Gottlob (University of Oxford, United Kingdom), Andreas Pieris (Brunel University, United Kingdom).

Data Extraction

566

Towards The Web of Concepts: Extracting Concepts from Large Datasets

Aditya Parameswaran (Stanford University, United States of America), Hector Garcia-Molina (Stanford University, United States of America), Anand Rajaraman (Kosmix Corporation, United States of America).

578

Exploiting Content Redundancy for Web Information Extraction

Pankaj Gulhane (Yahoo! Research Labs, Bangalore, India), Rajeev Rastogi (Yahoo! Research Labs, India), Srinivasan Sengamedu (Yahoo! Research Labs, Bangalore, India), Ashwin Tengli (Microsoft India Development Center, Bangalore, India).

588

Automatic Rule Refinement for Information Extraction

Bin Liu (University of Michigan, United States of America), Laura Chiticariu (IBM Almaden Research Center, United States of America), Vivian Chu (IBM Almaden Research Center, United States of America), H. Jagadish (University of Michigan, United States of America), Frederick Reiss (IBM Almaden Research Center, United States of America).

Privacy

598

Embellishing Text Search Queries To Protect User Privacy

HweeHwa Pang (Singapore Management University, Republic of Singapore), Xuhua Ding (Singapore Management University, Republic of Singapore), Xiaokui Xiao (Nanyang Technological University, Republic of Singapore).

608

Small Domain Randomization: Same Privacy, More Utility

Rhonda Chaytor (Simon Fraser University, Canada), Ke Wang (Simon Fraser University, Canada).

619

Nearest Neighbor Search with Strong Location Privacy

Stavros Papadopoulos (The Chinese University of Hong Kong, People's Republic of China), Spiridon Bakiras (City University of New York, United States of America), Dimitris Papadias (The Hong Kong University of Science and Technology, People's Republic of China).

Probabilistic and Uncertain Databases

630

UPI: A Primary Index for Uncertain Databases

Hideaki Kimura (Brown University, United States of America), Samuel Madden (Massachusetts Institute of Technology, United States of America), Stan Zdonik (Brown University, United States of America).

638

Ranking Continuous Probabilistic Datasets

Jian Li (University of Maryland College Park, United States of America), Amol Deshpande (University of Maryland College Park, United States of America).

650

Set Similarity Join on Probabilistic Data

Xiang Lian (The Hong Kong University of Science and Technology, People's Republic of China), Lei Chen (The Hong Kong University of Science and Technology, People's Republic of China).

Databases on Modern Hardware

660

Complex Event Detection at Wire Speed with FPGAs

Louis Woods (Eidgenössische Technische Hochschule Zürich, Switzerland), Jens Teubner (Eidgenössische Technische Hochschule Zürich, Switzerland), Gustavo Alonso (Eidgenössische Technische Hochschule Zürich, Switzerland).

670

Database Compression on Graphics Processors

Wenbin Fang (The Hong Kong University of Science and Technology, People's Republic of China), Bingsheng He (Nanyang Technological University, Republic of Singapore), Qiong Luo (The Hong Kong University of Science and Technology, People's Republic of China).

681

Aether: A Scalable Approach to Logging

Ryan Johnson (Carnegie Mellon University, United States of America), Ippokratis Pandis (Carnegie Mellon University, United States of America), Radu Stoica (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Manos Athanassoulis (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland).

Data Mining

693

Scalable Discovery of Best Clusters on Large Graphs

Kathy Macropol (University of California Santa Barbara, United States of America), Ambuj Singh (University of California Santa Barbara, United States of America).

703

An Architecture for Parallel Topic Models

Alexander Smola (Yahoo! Research Labs, United States of America), Shravan Narayana-murthy (Yahoo! Research Labs, India).

711

Keyword++: A Framework to Improve Keyword Search Over Entity Databases

Dong Xin (Microsoft Research, United States of America), Yeye He (University of Wisconsin-Madison, United States of America), Venkatesh Ganti (Google, United States of America).

Moving Object Databases

723

Swarm: Mining Relaxed Temporal Moving Object Clusters

Zhenhui Li (University of Illinois at Urbana-Champaign, United States of America), Bolin Ding (University of Illinois at Urbana-Champaign, United States of America), Jiawei Han (University of Illinois at Urbana-Champaign, United States of America), Roland Kays (New York State Museum, United States of America).

735

An Adaptive Updating Protocol for Reducing Moving Object Databases Workload

Su Chen (National University of Singapore, Republic of Singapore), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Zhenjie Zhang (National University of Singapore, Republic of Singapore).

747

Shortest Path Computation on Air Indexes

Georgios Kellaris (Singapore Management University, Republic of Singapore), Kyriakos Mouratidis (Singapore Management University, Republic of Singapore).

Probabilistic Data

758

Efficient and Effective Similarity Search over Probabilistic Data based on Earth Mover's Distance

Jia Xu (Northeastern University, People's Republic of China), Zhenjie Zhang (National University of Singapore, Republic of Singapore), Anthony Tung (National University of Singapore, Republic of Singapore), Ge Yu (Northeastern University, People's Republic of China).

770

Probabilistic XML via Markov Chains

Michael Benedikt (University of Oxford, United Kingdom), Evgeny Kharlamov (Free University of Bozen-Bolzano, Italy), Dan Olteanu (University of Oxford, United Kingdom), Pierre Senellart (Institut Télécom ParisTech, France).

Fuzzy, Probabilistic and Approximate Databases

782

MCDB-R: Risk Analysis in the Database

Peter Haas (IBM Research, United States of America), Christopher Jermaine (Rice University, United States of America), Subi Arumugam (University of Florida Gainesville, United States of America), Fei Xu (Microsoft Corporation, United States of America), Luis Perez (Rice University, United States of America), Ravi Jampani (University of Florida Gainesville, United States of America).

794

Scalable Probabilistic Databases with Factor Graphs and MCMC

Michael Wick (University of Massachusetts Amherst, United States of America), Andrew McCallum (University of Massachusetts Amherst, United States of America), Gerome Miklau (University of Massachusetts Amherst, United States of America).

Discovery and Exploration

805

On Multi-Column Foreign Key Discovery

Meihui Zhang (National University of Singapore, Republic of Singapore), Marios Hadjieleftheriou (AT&T Labs - Research, United States of America), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Cecilia Procopiuc (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

815

Explore or Exploit? Effective Strategies for Disambiguating Large Databases

Reynold Cheng (The University of Hong Kong, People's Republic of China), Eric Lo (Hong Kong Polytechnic University, People's Republic of China), Xuan Yang (The University of Hong Kong, People's Republic of China), Ming-Hay Luk (Hong Kong Polytechnic University, People's Republic of China), Xiang Li (The University of Hong Kong, People's Republic of China), Xike Xie (The University of Hong Kong, People's Republic of China).

Information Filtering and Dissemination

826

Building Ranked Mashups of Unstructured Sources with Uncertain Information

Mohamed Soliman (University of Waterloo, Canada), Ihab Ilyas (University of Waterloo, Canada), Mina Saleeb (University of Waterloo, Canada).

838

Computing Closed Skycubes

Chedy Raïssi (Institut National de Recherche en Informatique et Automatique, France), Jian Pei (Simon Fraser University, Canada), Thomas Kister (Laboratoire d'Informatique de Robotique et de Microélectronique de Montpellier).

Query Processing II

848

Generating Databases for Query Workloads

Eric Lo (Hong Kong Polytechnic University, People's Republic of China), Nick Cheng (Hong Kong Polytechnic University, People's Republic of China), Wing-Kai Hon (National Tsing Hua University, Taiwan).

860

Processing Top-k Join Queries

Minji Wu (Rutgers University, United States of America), Laure Berti-Équille (University of Rennes 1, France), Amélie Marian (Rutgers University, United States of America), Cecilia Procopiuc (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

871

Two-way Replacement Selection

Xavier Martinez-Palau (Universitat Politècnica de Catalunya, Spain), David Dominguez-Sal (Universitat Politècnica de Catalunya, Spain), Josep Lluís Larriba-Pey (Universitat Politècnica de Catalunya, Spain).

XML Data

882

XPath Whole Query Optimization

Sebastian Maneth (National Information and Communication Technologies Australia / University of New South Wales, Australia), Kim Nguyen (National Information and Communication Technologies Australia, Australia).

894

Fast Optimal Twig Joins

Nils Grimsmo (Norwegian University of Science and Technology, Norway), Truls Bjørklund (Norwegian University of Science and Technology, Norway), Magnus Hetland (Norwegian University of Science and Technology, Norway).

906

Destabilizers and Independence of XML Updates

Michael Benedikt (University of Oxford, United Kingdom), James Cheney (University of Edinburgh, United Kingdom).

Workflows, Transactions and Business Processes

918

Searching Workflows with Hierarchical Views

Ziyang Liu (Arizona State University, United States of America), Qihong Shao (Arizona State University, United States of America), Yi Chen (Arizona State University, United States of America).

928

Data-Oriented Transaction Execution

Ippokratis Pandis (Carnegie Mellon University, United States of America), Ryan Johnson (Carnegie Mellon University, United States of America), Nikos Hardavellas (Northwestern University, United States of America), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland).

940

Optimal Top-K Query Evaluation for Weighted Business Processes

Daniel Deutch (Tel Aviv University, Israel), Tova Milo (Tel Aviv University, Israel), Neoklis Polyzotis (University of California Santa Cruz, United States of America), Tom Yam (Tel Aviv University, Israel).

Scientific databases

952

Behavioral Simulations in MapReduce

Guozhang Wang (Cornell University, United States of America), Marcos Vaz Salles (Cornell University, United States of America), Benjamin Sowell (Cornell University, United States of America), Xun Wang (Cornell University, United States of America), Tuan Cao (Cornell University, United States of America), Alan Demers (Cornell University, United States of America), Johannes Gehrke (Cornell University, United States of America), Walker White (Cornell University, United States of America).

964

A-tree: A Structure for Storage and Modeling of Uncertain Multidimensional Arrays*

Tingjian Ge (University of Kentucky, United States of America), Stan Zdonik (Brown University, United States of America).

975

On Dense Pattern Mining in Graph Streams

Charu Aggarwal (IBM Thomas J. Watson Research Center, United States of America), Yao Li (University of Illinois at Chicago, United States of America), Philip Yu (University of Illinois at Chicago, United States of America), Ruoming Jin (Kent State University, United States of America).

Mobility and Spatial Queries

985

Efficient Proximity Detection among Mobile Users via Self-Tuning Policies

Man Lung Yiu (Hong Kong Polytechnic University, People's Republic of China), Leong Hou U (The University of Hong Kong, People's Republic of China), Simonas Šaltenis (Aalborg University, Denmark), Kostas Tzoumas (Aalborg University, Denmark).

997

k-Nearest Neighbors in Uncertain Graphs

Michalis Potamias (Boston University, United States of America), Francesco Bonchi (Yahoo! Research Labs, United States of America), Aristides Gionis (Yahoo! Research Labs, United States of America), George Kollios (Boston University, United States of America).

1009

Mining Significant Semantic Locations From GPS Data

Xin Cao (Nanyang Technological University, Republic of Singapore), Gao Cong (Nanyang

Technological University, Republic of Singapore), Christian Jensen (Aarhus University, Denmark).

Data Anonymization Techniques

1021

Boosting the Accuracy of Differentially Private Histograms Through Consistency

Michael Hay (University of Massachusetts Amherst, United States of America), Vibhor Rastogi (University of Washington, United States of America), Gerome Miklau (University of Massachusetts Amherst, United States of America), Dan Suciu (University of Washington, United States of America).

1033

rho-uncertainty: Inference-Proof Transaction Anonymization

Jianneng Cao (National University of Singapore, Republic of Singapore), Panagiotis Karras (National University of Singapore, Republic of Singapore), Chedy Raïssi (Institut National de Recherche en Informatique et Automatique, France), Kian-Lee Tan (National University of Singapore, Republic of Singapore).

1045

Minimizing Minimality and Maximizing Utility: Analyzing Method-based attacks on Anonymized Data

Graham Cormode (AT&T Labs - Research, United States of America), Ninghui Li (Purdue University, United States of America), Tiancheng Li (Purdue University, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

Querying and Integrating Probabilistic Databases

1057

Querying Probabilistic Information Extraction

Daisy Zhe Wang (University of California

Berkeley, United States of America), Michael Franklin (University of California Berkeley, United States of America), Minos Garofalakis (Technical University of Crete, Greece), Joseph Hellerstein (University of California Berkeley, United States of America).

1068

Read-Once Functions and Query Evaluation in Probabilistic Databases

Prithviraj Sen (Yahoo! Research Labs, United States of America), Amol Deshpande (University of Maryland College Park, United States of America), Lise Getoor (University of Maryland College Park, United States of America).

1080

Foundations of Uncertain-Data Integration

Parag Agrawal (Stanford University, United States of America), Anish Das Sarma (Yahoo! Research Labs, United States of America), Jeffrey Ullman (Stanford University, United States of America), Jennifer Widom (Stanford University, United States of America).

Database Design

1091

Identifying, Attributing and Describing Spatial Bursts

Michael Mathioudakis (University of Toronto, Canada), Nilesh Bansal (University of Toronto, Canada), Nick Koudas (University of Toronto, Canada).

1103

CORADD: Correlation Aware Database Designer for Materialized Views and Indexes

Hideaki Kimura (Brown University, United States of America), George Huo (Google, United States of America), Alexander Rasin (Brown University, United States of America), Samuel Madden (Massachusetts Institute

of Technology, United States of America), Stanley Zdonik (Brown University, United States of America).

1114

Regret-Minimizing Representative Databases

Danupon Nanongkai (Georgia Institute of Technology, United States of America), Atish Das Sarma (Georgia Institute of Technology, United States of America), Ashwin Lall (Georgia Institute of Technology, United States of America), Richard Lipton (Georgia Institute of Technology, United States of America), Jun Xu (Georgia Institute of Technology, United States of America).

Query Optimization

1125

An Access Cost-Aware Approach for Object Retrieval over Multiple Sources

Benjamin Arai (University of California Riverside, United States of America), Gautam Das (University of Texas at Arlington, United States of America), Dimitrios Gunopulos (University of Athens, Greece), Vagelis Hristidis (Florida International University, United States of America), Nick Koudas (University of Toronto, Canada).

1137

On the Stability of Plan Costs and the Costs of Plan Stability

M. Abhirama (Indian Institute of Science, India), Sourjya Bhaumik (Indian Institute of Science, India), Atreyee Dey (Indian Institute of Science, India), Harsh Shrimal (Indian Institute of Science, India), Jayant Haritsa (Indian Institute of Science, India).

1149

Xplus: A SQL-Tuning-Aware Query Optimizer

Herodotos Herodotou (Duke University, United States of America), Shivnath Babu (Duke University, United States of America).

Graph and Pattern Matching

1161

Graph Homomorphism Revisited for Graph Matching

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (Harbin Institute of Technology, People's Republic of China), Shuai Ma (University of Edinburgh, United Kingdom), Hongzhi Wang (Harbin Institute of Technology, People's Republic of China), Yinghui Wu (University of Edinburgh, United Kingdom).

1173

SigMatch: Fast and Scalable Multi-Pattern Matching

Ramakrishnan Kandhan (University of Wisconsin-Madison, United States of America), Nikhil Teletia (University of Wisconsin-Madison, United States of America), Jignesh Patel (University of Wisconsin-Madison, United States of America).

1185

SAPPER: Subgraph Indexing and Approximate Matching in Large Graphs

Shijie Zhang (Case Western Reserve University, United States of America), Jiong Yang (Case Western Reserve University, United States of America), Wei Jin (Case Western Reserve University, United States of America).

Indexing Techniques

1195

Tree Indexing on Solid State Drives

Yinan Li (University of Wisconsin-Madison, United States of America), Bingsheng He (The Hong Kong University of Science and Technology, People's Republic of China), Robin Jun Yang (The Hong Kong University of Science and Technology, People's Republic of China), Qiong Luo (The Hong Kong University of Science and Technology, People's Republic of China), Ke

Yi (The Hong Kong University of Science and Technology, People's Republic of China).

1207

Efficient B-tree Based Indexing for Cloud Data Processing

Sai Wu (National University of Singapore, Republic of Singapore), Dawei Jiang (National University of Singapore, Republic of Singapore), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Kun-Lung Wu (IBM Thomas J. Watson Research Center, United States of America).

1219

Trie-Join: Efficient Trie-based String Similarity Joins with Edit-Distance Constraints

Jiannan Wang (Tsinghua University, People's Republic of China), Guoliang Li (Tsinghua University, People's Republic of China), Jianhua Feng (Tsinghua University, People's Republic of China).

Query Processing III

1231

VoR-Tree: R-trees with Voronoi Diagrams for Efficient Processing of Spatial Nearest Neighbor Queries

Mehdi Sharifzadeh (Google, United States of America), Cyrus Shahabi (University of Southern California, United States of America).

1243

Efficient RkNN Retrieval with Arbitrary Non-Metric Similarity Measures

Deepak P (IBM India Research Lab, India), Prasad Deshpande (IBM India Research Lab, India).

1255

Efficient Skyline Evaluation over Partially Ordered Domains

Shiming Zhang (The University of Hong Kong, People's Republic of China), Nikos Mamoulis (The University of Hong Kong, People's Republic of China), Ben Kao (The University of Hong Kong, People's Republic of China), David Cheung (The University of Hong Kong, People's Republic of China).

Streaming and Sensor Data

1267

Achieving High Output Quality under Limited Resources through Structure-based Spilling in XML Streams

Mingzhu Wei (Worcester Polytechnic Institute, United States of America), Elke Rundensteiner (Worcester Polytechnic Institute, United States of America), Murali Mani (University of Michigan at Flint, United States of America).

1279

Dynamic Join Optimization in Multi-Hop Wireless Sensor Networks

Svilen Mihaylov (University of Pennsylvania, United States of America), Marie Jacob (University of Pennsylvania, United States of America), Zachary Ives (University of Pennsylvania, United States of America), Sudipto Guha (University of Pennsylvania, United States of America).

1291

Database-support for Continuous Prediction Queries over Streaming Data

Mert Akdere (Brown University, United States of America), Uğur Çetintemel (Brown University, United States of America), Eli Upfal (Brown University, United States of America).

1302

Conditioning and Aggregating Uncertain Data Streams: Going Beyond Expectations

Thanh Tran (University of Massachusetts Amherst, United States of America), Andrew McGregor (University of Massachusetts

Amherst, United States of America), Yanlei Diao (University of Massachusetts Amherst, United States of America), Liping Peng (University of Massachusetts Amherst, United States of America), Anna Liu (University of Massachusetts Amherst, United States of America).

Information Integration and Retrieval

1314

TRAMP: Understanding the Behavior of Schema Mappings through Provenance

Boris Glavic (University of Zurich, Switzerland), Gustavo Alonso (Eidgenössische Technische Hochschule Zürich, Switzerland), Renée Miller (University of Toronto, Canada), Laura Haas (IBM Research, United States of America).

1326

Entity Resolution with Evolving Rules

Steven Whang (Stanford University, United States of America), Hector Garcia-Molina (Stanford University, United States of America).

1338

Annotating and Searching Web Tables Using Entities, Types and Relationships

Girija Limaye (Indian Institute of Technology - Bombay, India), Sunita Sarawagi (Indian Institute of Technology - Bombay, India), Soumen Chakrabarti (Indian Institute of Technology - Bombay, India).

Data Mining, Copy Detection and Data Publication

1348

Interesting-Phrase Mining for Ad-Hoc Text Analytics

Srikanta Bedathur (Max-Planck Institute for

Informatics, Germany), Klaus Berberich (Max-Planck Institute for Informatics, Germany), Jens Dittrich (Saarland University, Germany), Nikos Mamoulis (Max-Planck Institute for Informatics, Germany), Gerhard Weikum (Max-Planck Institute for Informatics, Germany).

1358

Global Detection of Complex Copying Relationships Between Sources

Xin Dong (AT&T Labs - Research, United States of America), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

1370

Fragments and Loose Associations: Respecting Privacy in Data Publishing

Sabrina De Capitani di Vimercati (Università degli Studi di Milano, Italy), Sara Foresti (Università degli Studi di Milano, Italy), Sushil Jajodia (George Mason University Fairfax, United States of America), Stefano Paraboschi (University of Bergamo, Italy), Pierangela Samarati (Università degli Studi di Milano, Italy).

Industrial Applications and Experience Sessions

Real-Time and Stream Processing

1382

Net-Fli: On-the-fly Compression, Archiving and Indexing of Streaming Network Traffic

Francesco Fusco (IBM Research - Zurich, Switzerland), Marc Stoecklin (IBM Research - Zurich, Switzerland), Michael Vlachos (IBM Research - Zurich, Switzerland).

1394

From a Stream of Relational Queries to Distributed Stream Processing

Qiong Zou (IBM China Research Lab, People's Republic of China), Huayong Wang (IBM China Research Lab, People's Republic of China), Robert Soulé (New York University, United States of America), Martin Hirzel (IBM Thomas J. Watson Research Center, United States of America), Henrique Andrade (IBM Thomas J. Watson Research Center, United States of America), Buğra Gedik (IBM Thomas J. Watson Research Center, United States of America), Kun-Lung Wu (IBM Thomas J. Watson Research Center, United States of America).

1406

UASMAs (Universal Automated SNP Mapping Algorithms): a set of algorithms to instantaneously map SNPs in real time to aid functional SNP discovery

James Mah (Institute for Infocomm Research, Republic of Singapore), Danny Poo (National University of Singapore, Republic of Singapore), Shaojiang Cai (National University of Singapore, Republic of Singapore).

Advances in Database Systems

1414

FlashStore: High Throughput Persistent Key-Value Store

Biplob Debnath (University of Minnesota, United States of America), Sudipta Sengupta (Microsoft Research - Redmond, United States of America), Jin Li (Microsoft Research - Redmond, United States of America).

1426

MEET DB2: Automated Database Migration Evaluation

Reynold Xin (University of California Berkeley, United States of America), Patrick Dantsangale (IBM Research, United Kingdom), Sam Lightstone (IBM Research, Canada), William McLaren (IBM, United States of America), Steve Schormann (IBM Research, Canada), Maria Schwenger (IBM, United States of America).

1435

SSD Bufferpool Extensions for Database Systems

Mustafa Canim (University of Texas at Dallas, United States of America), George Mihaila (IBM Thomas J. Watson Research Center, United States of America), Bishwaranjan Bhattacharjee (IBM Thomas J. Watson Research Center, United States of America), Kenneth Ross (Columbia University, United States of America), Christian Lang (IBM Thomas J. Watson Research Center, United States of America).

Data Warehousing and Mining

1447

DataGarage: Warehousing Massive Performance Data on Commodity Servers

Charles Lobo (Microsoft Corporation, United States of America), Slawek Smyl (Microsoft Corporation, United States of America), Suman Nath (Microsoft Research, United States of America).

1459

Cheetah: A High Performance, Custom Data Warehouse on Top of MapReduce

Songting Chen (Turn Incorporated, United States of America).

1469

Distance-Based Outlier Detection: Consolidation and Renewed Bearing

Gustavo Orair (Universidade Federal de Minas Gerais, Brazil), Carlos Teixeira (Universidade Federal de Minas Gerais, Brazil), Ye Wang (The Ohio State University, United States of America), Wagner Meira (Universidade Federal de Minas Gerais, Brazil), Srinivasan Parthasarathy (The Ohio State University, United States of America).

Mobile and Social Applications

1481

Adaptive Logging for Mobile Device

Young-Seok Kim (Samsung Advanced Institute of Technology, Republic of Korea), Heegyu Jin (Samsung Advanced Institute of Technology, Republic of Korea), Kyoung-Gu Woo (Samsung Advanced Institute of Technology, Republic of Korea).

1493

RoadTrack: Scaling Location Updates for Mobile Clients on Road Networks with Query Awareness

Peter Pesti (Georgia Institute of Technology, United States of America), Ling Liu (Georgia Institute of Technology, United States of America), Bhuvan Bamba (Georgia Institute of Technology, United States of America), Arun Iyengar (IBM Thomas J. Watson Research Center, United States of America), Matt Weber (Georgia Institute of Technology, United States of America).

1505

Confucius and Its Intelligent Disciples: Integrating Social with Search

Xiance Si (Google Research, People's Republic of China), Edward Chang (Google Research, People's Republic of China), Zoltán Gyöngyi (Google Research, United States of America), Maosong Sun (Tsinghua University, People's Republic of China).

Demonstrations

Core Database Technologies (12 Demos)

1517

The Picasso Database Query Optimizer Visualizer

Jayant Haritsa (Indian Institute of Science, India).

1521

CODS: Evolving Data Efficiently and Scalably in Column Oriented Databases

Ziyang Liu (Arizona State University, United States of America), Sivaramakrishnan Natarajan (Arizona State University, United States of America), Bin He (IBM Almaden Research Center, United States of America), Hui-I Hsiao (IBM Almaden Research Center, United States of America), Yi Chen (Arizona State University, United States of America).

1525

Efficient Event Processing through Reconfigurable Hardware for Algorithmic Trading

Mohammad Sadoghi (University of Toronto, Canada), Hans-Arno Jacobsen (University of Toronto, Canada), Martin Labrecque (University of Toronto, Canada), Warren Shum (University of Toronto, Canada), Harsh Singh (University of Toronto, Canada).

1529

CareDB: A Context and Preference-Aware Location-Based Database System

Justin Levandoski (University of Minnesota, United States of America), Mohamed Mokbel (University of Minnesota, United States of America), Mohamed Khalefa (University of Minnesota, United States of America).

1533

Cloudy: A Modular Cloud Storage System

Donald Kossmann (Eidgenössische Technische Hochschule Zürich, Switzerland), Tim Kraska (University of California Berkeley, United States of America), Simon Loesing (Eidgenössische Technische Hochschule Zürich, Switzerland), Stephan Merkli (Eidgenössische Technische Hochschule Zürich, Switzerland), Raman Mittal (Eidgenössische Technische Hochschule Zürich, Switzerland), Flavio Pfaffhauser (Eidgenössische Technische Hochschule Zürich, Switzerland).

1537

Geospatial Stream Query Processing using Microsoft SQL Server StreamInsight

Seyed Jalal Kazemitabar (University of Southern California, United States of America), Ugur Demiryurek (University of Southern California, United States of America), Mohamed Ali (Microsoft Corporation, United States of America), Afsin Akdogan (University of Southern California, United States of America), Cyrus Shahabi (University of Southern California, United States of America).

1541

Using XMorph to Transform XML Data

Curtis Dyreson (Utah State University, United States of America), Sourav Bhowmick (Nanyang Technological University, Republic of Singapore), Kirankanth Mallampalli (Utah State University, United States of America).

1545

Active Complex Event Processing: Applications in Real-Time Health Care

Di Wang (Worcester Polytechnic Institute, United States of America), Elke Rundensteiner (Worcester Polytechnic Institute, United States of America), Richard Ellison (University of Massachusetts Medical School, United States of America), Han Wang (Worcester Polytechnic Institute, United States of America).

1549

Thirteen New Players in the Team: A Ferry-based LINQ to SQL Provider

Tom Schreiber (Universität Tübingen, Germany), Simone Bonetti (Universität Tübingen, Germany), Torsten Grust (Universität Tübingen, Germany), Manuel Mayr (Universität Tübingen, Germany), Jan Rittinger (Universität Tübingen, Germany).

1553

AXART - Enabling Collaborative Work with AXML Artifacts

Bogdan Marinoiu (SAP Labs, France), Serge Abiteboul (Institut National de Recherche en Informatique et en Automatique, France), Pierre Bourhis (Institut National de Recherche en Informatique et en Automatique, France), Alban Galland (Institut National de Recherche en Informatique et en Automatique, France).

1557

iFlow: An Approach for Fast and Reliable Internet-Scale Stream Processing Utilizing Detouring and Replication

Christopher McConnell (State University of New York University at Albany, United States of America), Fan Ping (State University of New York University at Albany, United States of America), Jeong-Hyon Hwang (State University of New York University at Albany, United States of America).

1561

Peer coordination through distributed triggers

Verena Kantere (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Maher Manoubi (University of Ottawa, Canada), Iluju Kiringa (University of Ottawa, Canada), Timos Sellis (National Technical University of Athens, Greece), John Mylopoulos (University of Toronto, Canada).

Web Data Management and Information Retrieval Systems (12 Demos)

1565

Seaform: Search-As-You-Type in Forms

Hao Wu (Tsinghua University, People's Republic of China), Guoliang Li (Tsinghua University, People's Republic of China), Chen Li (University of California Irvine, United States of America), Lizhu Zhou (Tsinghua University, People's Republic of China).

1569

TimeTrails: A System for Exploring Spatio-Temporal Information in Documents

Jannik Strötgen (University of Heidelberg, Germany), Michael Gertz (University of Heidelberg, Germany).

1573

QUICK: Expressive and Flexible Search over Knowledge Bases and Text Collections

Jeffrey Pound (University of Waterloo, Canada), Ihab Ilyas (University of Waterloo, Canada), Grant Weddell (University of Waterloo, Canada).

1577

Transforming XML Documents as Schemas Evolve

Jarek Gryz (York University, Canada), Marcin Kwietniewski (York University, Canada), Stephanie Hazlewood (IBM, Canada), Paul Van Run (IBM, Canada).

1581

XSACT: A Comparison Tool for Structured Search Results

Ziyang Liu (Arizona State University, United States of America), Sivaramakrishnan Nataraajan (Arizona State University, United States of America), Peng Sun (Arizona State University, United States of America), Stephen Boohar (Arizona State University, United States of America), Tim Meehan (Arizona State University, United States of America), Robert Winkler (Arizona State University, United States of America), Yi Chen (Arizona State University, United States of America).

1585

ObjectRunner: Lightweight, Targeted Extraction and Querying of Structured Web Data

Talel Abdessalem (Télécom ParisTech, France),

Bogdan Cautis (Télécom ParisTech, France), Nora Derouiche (Télécom ParisTech, France).

1589

ROXXI: Reviving witness dOcuments to eX-plore eXtracted Information

Shady Elbassuoni (Max-Planck Institute for Informatics, Germany), Katja Hose (Max-Planck Institute for Informatics, Germany), Steffen Metzger (Max-Planck Institute for Informatics, Germany), Ralf Schenkel (Max-Planck Institute for Informatics, Germany).

1593

EXTRUCT: Using Deep Structural Information in XML Keyword Search

Arash Termehchy (University of Illinois at Urbana-Champaign, United States of America), Marianne Winslett (University of Illinois at Urbana-Champaign, United States of America).

1597

SQL QueRIE Recommendations

Javad Akbarnejad (San Jose State University, United States of America), Gloria Chatzopoulou (University of California Riverside, United States of America), Magdalini Eirinaki (San Jose State University, United States of America), Suju Koshy (San Jose State University, United States of America), Sarika Mittal (San Jose State University, United States of America), Duc On (San Jose State University, United States of America), Neoklis Polyzotis (University of California Santa Cruz, United States of America), Jothi Swarubini Vindhiya Varman (San Jose State University, United States of America).

1601

P2PDocTagger: Content management through automated P2P collaborative tagging

Hock Hee Ang (Nanyang Technological University, Republic of Singapore), Vivekanand

Gopalkrishnan (Nanyang Technological University, Republic of Singapore), Wee Keong Ng (Nanyang Technological University, Republic of Singapore), Steven C.H. Hoi (Nanyang Technological University, Republic of Singapore).

1605

InZeit: Efficiently Identifying Insightful Time Points

Vinay Setty (Max-Planck Institute for Informatics, Germany), Srikanta Bedathur (Max-Planck Institute for Informatics, Germany), Klaus Berberich (Max-Planck Institute for Informatics, Germany), Gerhard Weikum (Max-Planck Institute for Informatics, Germany).

1609

iAVATAR: An Interactive Tool for Finding and Visualizing Visual-Representative Tags in Image Search

Aixin Sun (Nanyang Technological University, Republic of Singapore), Sourav Bhowmick (Nanyang Technological University, Republic of Singapore), Yao Liu (Nanyang Technological University, Republic of Singapore).

Data Extraction, Integration and Mining (8 Demos)

1613

Deep Web Integration with VisQI

Thomas Kabisch (Humboldt Universität zu Berlin, Germany), Eduard Dragut (University of Illinois at Chicago, United States of America), Clement Yu (University of Illinois at Chicago, United States of America), Ulf Leser (Humboldt Universität zu Berlin, Germany).

1617

SOLOMON: Seeking the Truth Via Copying Detection

Xin Dong (AT&T Labs - Research, United States of America), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T

Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

1621

Just-in-time Data Integration in Action

Martin Hentschel (Eidgenössische Technische Hochschule Zürich, Switzerland), Laura Haas (IBM Almaden Research, United States of America), Renée Miller (University of Toronto, Canada).

1625

Massively Parallel Data Analysis with PACTs on Nephele

Alexander Alexandrov (Technische Universität Berlin, Germany), Dominic Battré (Technische Universität Berlin, Germany), Stephan Ewen (Technische Universität Berlin, Germany), Max Heimes (Technische Universität Berlin, Germany), Fabian Hueske (Technische Universität Berlin, Germany), Odej Kao (Technische Universität Berlin, Germany), Volker Markl (Technische Universität Berlin, Germany), Erik Nijkamp (Technische Universität Berlin, Germany), Daniel Warneke (Technische Universität Berlin, Germany).

1629

Using Sentinel Technology in the TARGIT BI Suite

Morten Middelfart (TARGIT A/S, Denmark), Torben Bach Pedersen (Aalborg University, Denmark).

1633

CoDA: Interactive Cluster Based Concept Discovery

Stephan Günemann (RWTH Aachen University, Germany), Ines Färber (RWTH Aachen University, Germany), Hardy Kremer (RWTH Aachen University, Germany), Thomas Seidl (RWTH Aachen University, Germany).

1637

Keymantic: Semantic Keyword-based Searching in Data Integration Systems

Sonia Bergamaschi (University of Modena and Reggio Emilia, Italy), Elton Domnori (University of Modena and Reggio Emilia, Italy), Francesco Guerra (University of Modena and Reggio Emilia, Italy), Mirko Orsini (University of Modena and Reggio Emilia, Italy), Raquel Trillo Lado (University of Zaragoza, Spain), Yannis Velegarakis (University of Trento, Italy).

1641

Data Auditor: Exploring Data Quality and Semantics using Pattern Tableau

Lukasz Golab (AT&T Labs - Research, United States of America), Howard Karloff (AT&T Labs - Research, United States of America), Flip Korn (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

Tutorials

1645

Distributed Caching Platforms

Anil Nori (Microsoft Corporation, United States of America).

1647

Big Data and Cloud Computing: New Wine or just New Bottles?

Divyakant Agrawal (University of California Santa Barbara, United States of America), Sudipto Das (University of California Santa Barbara, United States of America), Amr El Abbadi (University of California Santa Barbara, United States of America).

1649

Techniques for Similarity Searching in Multimedia Databases

Hanan Samet (University of Maryland College Park, United States of America).

1651

Event Processing – past, present and future

Opher Etzion (IBM Haifa Research Lab, Israel).

1653

Similarity Search and Mining in Uncertain Databases

Matthias Renz (Ludwig Maximilian University of Munich, Germany), Reynold Cheng (The University of Hong Kong, People's Republic of China), Hans-Peter Kriegel (Ludwig Maximilian University of Munich, Germany).

1655

Data Management and Mining in Internet Ad Systems

S. Muthukrishnan (Rutgers U., and Google Inc., United States of America).

Panels

Keynote Sessions

1657

Cloud Databases: What's New?

Daniel Abadi (Yale University, United States of America), Michael Carey (University of California Irvine, United States of America), Surajit Chaudhuri (Microsoft Research, United States of America), Hector Garcia-Molina (Stanford University, United States of America), Jignesh Patel (University of Wisconsin-Madison, United States of America).

1658

Time for Our Field to Grow Up

Anastassia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Laura Haas (IBM Almaden Research Center, United

States of America), HV Jagadish (University of Michigan, United States of America), David

Maier (Portland State University, United States of America), Tamer Özsu (University of Waterloo, Canada).