

SSDBM 2010 Program

Wednesday, June 30th, 2010

8:00 Registration

8:45 Opening

9:00 **Keynote talk 1:** Daniel J. Abadi. *Tradeoffs Between Parallel Database Systems, Hadoop, and HadoopDB as Platforms for Petabyte-Scale Analysis*

10:00 Coffee break

10:30 **Session 1: Query Processing**

Markus Bestehorn, Klemens Böhm, Patrick Bradley, Erik Buchmann. *Deriving Spatio-Temporal Query Results in Sensor Networks*

George Valkanas, Apostolos N. Papadopoulos. *Efficient and Adaptive Distributed Skyline Computation*

Romans Kasperovics, Michael H. Böhlen, Johann Gamper. *On the Efficient Construction of Multislices from Recurrences*

Mario A. Nascimento, Romulo A.E. Alencar, Angelo Brayner. *Optimizing Query Processing in Cache-Aware Wireless Sensor Networks*

Katja Seidler, Eric Peukert, Gregor Hackenbroich, Wolfgang Lehner. *Approximate Query Answering and Result Refinement on XML Data* (short paper)

Ariel Cary, Ouri Wolfson, Naphtali Rische. *Efficient and Scalable Method for Processing Top-k Spatial Boolean Queries* (short paper)

12:30 Lunch break

14:00 **Session 2: Scientific Data Management and Analysis**

Shen-Shyang Ho, Wenqing Tang, W. Timothy Liu, Markus Schneider. *A Framework for Moving Sensor Data Query and Retrieval of Dynamic Atmospheric Events*

Keith Grochow, Bill Howe, Ed Lazowska, Roger Barga, Mark Stoermer. *Client + Cloud: Evaluating Seamless Architectures for Visual Data Analytics in the Ocean Sciences*

YongChul Kwon, Dylan Nunley, Jeffrey P. Gardner, Magdalena Balazinska, Bill Howe, Sarah Loebman. *Scalable Clustering Algorithm for N-body Simulations in a Shared-nothing Cluster*

Viswanath Nandigam, Chaitan Baru, Christopher Crosby. *Database Design for High-Resolution LIDAR Topography Data* (short paper)

Andrei Aiordăchioaie, Peter Baumann. *PetaScope: An Open-source Implementation of the OGC WCS Geo Service Standards Suite* (demo paper)

Hans-Peter Kriegel, Peer Kröger, Christiaan Hendrikus van der Meijden, Henriette Obermaier, Joris Peters, Matthias Renz. *Towards Archaeo-Informatics: Scientific Data Management for Archaeobiology* (short paper)

15:30 Coffee break

16:00 **Session 3: Data Mining**

Shirong Li, Shijie Zhang, Jiong Yang. *DESSIN: Mining Dense Subgraph Patterns in a Single Graph*

Arnab Bhattacharya, Abhishek Bhowmick, Ambuj K. Singh. *Finding Top-k Similar Pairs of Objects Annotated with Terms from an Ontology*

Hua Lu, Linhao Xu. *Identifying the Most Influential User Preference from an Assorted Collection*

Philipp Kranen, Stephan Günnemann, Sergej Fries, Thomas Seidl. *MC-Tree: Improving Bayesian Anytime Classification*

Zaigham Faraz Siddiqui, Myra Spiliopoulou. *Tree Induction over Perennial Objects*

Mark Brightwell, Anastasia Ailamaki, Anna Suwalska. *Non-intrusive Quality Analysis of Monitoring Data* (short paper)

Martin Hahmann, Dirk Habich, Wolfgang Lehner. *Visual Decision Support for Ensemble-Clustering* (demo paper)

18:00 Reception – Presentation of Demos and Posters (Location: Studio Villa Bosch)

Thursday, July 1st, 2010

9:00 **Keynote talk 2:** Roger S. Barga. *Emerging Trends and Converging Technologies in Data Intensive Scalable Computing*

10:00 Coffee break

10:30 **Session 4: Indexes and Data Representation**

Zeyar Aung, See-Kiong Ng. *An Indexing Scheme for Fast and Accurate Chemical Fingerprint Database Searching*

Julio López, Leonardo Ramírez-Guzmán, Jacobo Bielak, David O'Hallaron. *BEMC: A Searchable, Compressed Representation for Large Seismic Wavefields*

Ekow Otoo, Doron Rotem, Shih-Chiang Tsao. *Dynamic Data Reorganization for Energy Savings in Disk Storage Systems*

Eric Perlman, Randal Burns, Michael Kazhdan, Rebecca R. Murphy, William P. Ball, Nina Amenta. *Organization of Data in Non-Convex Spatial Domains*

Gaoping Zhu, Xuemin Lin, Wenjie Zhang, Wei Wang, Haichuan Shang. *PrefIndex: An Efficient Supergraph Containment Search Technique*

Jianting Zhang, Simin You. *Supporting Web-based Visual Exploration of Large-Scale Raster Geospatial Data Using Binned Min-Max Quadtree*

12:30 Lunch break

14:00 **Session 5: Scientific Workflow and Provenance**

David Koop, Emanuele Santos, Bela Bauer, Matthias Troyer, Juliana Freire, Cláudio T. Silva. *Bridging Workflow and Data Provenance using Strong Links*

Anish Das Sarma, Martin Theobald, Jennifer Widom. *LIVE: A Lineage-Supported Versioned DBMS*

Christoph Langguth, Heiko Schuldt. *Optimizing Resource Allocation for Scientific Workflows using Advance Reservations*

Pierre Mouallem, Daniel Crawl, Ilkay Altintas, Mladen Vouk, Ustun Yildiz. *A Fault-Tolerance Architecture for Kepler-based Distributed Scientific Workflows* (short paper)

Satya S. Sahoo, Olivier Bodenreider, Pascal Hitzler, Amit Sheth, Krishnaprasad Thirunarayan. *Provenance Context Entity (PaCE): Scalable Provenance Tracking for Scientific RDF Data* (short paper)

Paolo Missier, Stian Soiland-Reyes, Stuart Owen, Wei Tan, Alexandra Nenadic, Ian Dunlop, Alan Williams, Tom Oinn, Carole Goble. *Taverna, reloaded* (short paper)

15:30 Coffee break

16:00 **Session 6: Similarity**

Michael E. Houle, Hans-Peter Kriegel, Peer Kröger, Erich Schubert, Arthur Zimek. *Can Shared-Neighbor Distances Defeat the Curse of Dimensionality?*

Tobias Emrich, Franz Graf, Hans-Peter Kriegel, Matthias Schubert, Marisa Thoma. *Optimizing All-Nearest-Neighbor Queries with Trigonometric Pruning*

Astrid Rheinländer, Martin Knobloch, Nicky Hochmuth, Ulf Leser. *Prefix Tree Indexing for Similarity Search and Similarity Joins on Genomic Data*

Matthias Schubert, Tobias Emrich, Hans-Peter Kriegel, Marisa Thoma, Franz Graf. *Similarity Estimation using Bayes Ensembles*

Thomas Bernecker, Tobias Emrich, Franz Graf, Hans-Peter Kriegel, Peer Kröger, Matthias Renz, Erich Schubert, Arthur Zimek. *Subspace Similarity Search: Efficient k -NN Queries in Arbitrary Subspaces* (short paper)

19:00 Conference Banquet (Location: NH Hotels, Bergheimer Strasse 91)

Friday, July 2nd, 2010

9:00 Session 7: Data Stream Processing

Hua Lu, Yongluan Zhou, Jonas Haustad. *Continuous Skyline Monitoring over Distributed Data Streams*

Htoo Htet Aung, Kian-Lee Tan. *Discovery of Evolving Convoys*

Michael Daum, Frank Lauterwald, Philipp Baumgärtel, Klaus Meyer-Wegener. *Propagation of Densities of Streaming Data within Query Graphs*

Mingyan Gao, Xiaoyan Yang, Ramesh Jain, Beng Chin Ooi. *Spatio-Temporal Event Stream Processing in Multimedia Communication Systems*

Mohammed Al-Kateb, Byung Suk Lee. *Stratified Reservoir Sampling over Heterogeneous Data Streams*

11:00 Coffee break

11:30 Panel “MapReduce and DBMS systems: Complementary or Competitive Technologies?” Panel members: Daniel Abadi (Yale University), Roger Barga (Microsoft Research), Julio Lopez (Carnegie Mellon University), Norman May (SAP Research)

12:30 Conference closing