VLDB SINGAPORE 2010

36th International Conference on Very Large Data Bases

Singapore : 13 to 17 Sept 2010,
Grand Copthorne Waterfront Hotel

NUS National University of Singapore

SMU SINGAPORE MANAGEMENT UNIVERSITY

NANYANG TECHNOLOGICAL UNIVERSITY

Very Large Data Bases

VLDB 2010

September 13th – September 17th, 2010
Grand Copthorne Waterfront Hotel

http://www.vldb2010.org

TABLE OF CONTENTS

Created By: Sourav S Bhowmick
In Cooperation With: VLDB Organizing Committee,
Singapore Tourism Board, Mr Lai Chee Keong (NTU)

WELCOME MESSAGE FROM GENERAL CHAIRS

On behalf of the conference committee for VLDB 2010, it is our pleasure to welcome you to Singapore for the 36th International Conference on Very Large Data Bases. This is the second time that this conference is being held in Singapore; the first time was 26 years ago in 1984. The conference is jointly organized by the National University of Singapore, Nanyang Technological University, Singapore Management University, and the VLDB Endowment.

VLDB is a premier annual international forum for data management and database researchers, vendors, practitioners, application developers, and users. The conference will feature research talks, industry presentations, panels, tutorials, demonstrations, and 11 workshops. It will cover current issues in data management, database and information systems research. Data management and databases remain among the main technological cornerstones of emerging applications of the twenty-first century.

Singapore is situated at the southern tip of the Malaysian Peninsula in South-East Asia. It is a small and prosperous cosmopolitan nation that offers visitors the cultural diversity of its four main ethnic groups, namely, Chinese, Malays, Indians and Eurasians. Singapore is a vibrant Science & Technology hub where several universities and numerous public and corporate research institutes perform high-quality research over a broad range of disciplinary and cross-disciplinary areas. Singapore benefits from a knowledge-based economy with a national R&D spending target of at least 3% of GDP by 2010. With museums, theme parks, beaches, nature parks and two new integrated resorts as well as bustling shopping and dining havens, Singapore is an excellent place to unwind.

The conference would not have been possible if not for the efforts of many people. Thanks are due to the Technical Program Chair - Dr. Elisa BERTINO, and her team of track program chairs – Dr. Kian Lee TAN, Dr. Paolo ATZENI, Dr. Philip S. YU, and Dr. Masaru KITSUREGAWA and their PC members for producing an exciting technical programme. Thanks also to the efforts of the workshop co-chairs – Dr. Amol DESHPANDE, Dr. Zachary G. IVES, and Dr. Anthony Kum Hoe TUNG, workshops' organizers and their committee members, the VLDB Endowment Liaison – Dr. David LOMET, the Local Arrangement Co-Chairs – Dr. Sourav S BHOWMICK, and Dr. Mong Li LEE, the Sponsorship Co-Chairs – Dr. Dan SUCIU and Dr. Limsoon WONG, the Publicity co-chairs – Dr. Stéphane BRESSAN, Dr. Umeshwar DAYAL, and Dr. Christian S. JENSEN. the proceedings co-chairs – Dr. Y.C. TAY and Dr. Yi CHEN, the Registration Chair – Dr. Chee Yong CHAN, the Treasurer – Dr. Wynne HSU, webmaster – Mr. Zaini MOHD, the Conference Secretary – Ms. Siew Foong HO, and other committee chairs, co-chairs, and members.

We are grateful for the generous support of our sponsors, Platinum Sponsors - Microsoft and Oracle; Gold Sponsors - Google, IBM, SAP, and Sybase; Silver Sponsors - ADSC, Yahoo!, TPC, and HP; Supporter - Sun Microsystems. We also wish to thank the following supporting organizations - A*Star and Singapore Tourism Board; supporting societies - Korean Institute of Information Scientists and Engineers, COMAD (India), The Database Society of Japan, and the China Computer Federation Database Technical Society.

Last but not least, our sincere thanks go to the authors of the papers, the speakers, and all the participants of VLDB 2010 who have made this conference a resounding success.

Welcome and enjoy the conference and have a good time in Singapore!

Tok Wang LING, National University of Singapore, Singapore
Krithi RAMAMRITHAM, Indian Institute of Technology Bombay, India
VLDB 2010 General Chairs

WELCOME MESSAGE FROM PROGRAM CHAIRS

A warm welcome to the 36th International Conference on Very Large Data Bases (VLDB) in Singapore. The annual VLDB Conference is the premier international venue for the dissemination and exchange of research advances and ideas in the broad area of data management, both in the development and application of new technology. We hope you find both the conference and its location enjoyable.

This year's edition of VLDB received 686 submissions overall. For the research tracks: The Core Database Technology track received 280 submissions, of these 19 were rejected without review because of formatting violations and 48 were accepted. The acceptance rate computed with respect to the reviewed papers is 18.4%.

The Infrastructure for Information Systems track received 215 submissions, of these 11 were rejected without review because of formatting violations and 33 were accepted. The acceptance rate computed with respect to the reviewed papers is 16.1%.

In addition to the traditional program committees, papers were also reviewed through a "journal track". 119 submissions were received in this track during this year, of which 30 were accepted. In addition, there were 6 acceptances from other submission periods.

In addition, the Industrial Applications and Experience track received 51 submissions, of which 12 (23.5 %) were accepted.

The Demonstrations track received 89 submissions, of which 32 (36 %) were accepted.

The conference program also includes two keynote talks. We are very fortunate that two outstanding researchers have accepted our invitation to serve as keynote speakers. Divesh Srivastava will give the first keynote talk, focusing on real time data analysis, whereas Paul Matsudaira will give the second keynote talk, focusing on the challenges of managing very large 3D+ datasets resulting from high-end biological images.

The conference is complemented by a strong program of 11 workshops selected by Amol Deshpande,University of Maryland, Zachary G. Ives, University of Pennsylvania, and Anthony Kum Hoe Tung, National University of Singapore, Singapore. Of these, the VLDB PhD workshop continues to be encouraged and embraced by the VLDB Program Chairs. It was chaired this year by John Mylopoulos, University of Toronto, Lizhu Zhou, Tsinghua University, and Xiaofang Zhou, University of Queensland.

The Tutorial Program Chairs, Guozhu Dong, Wright State University, and Hwee Hwa Pang, Singapore Management University, selected 6 tutorials covering a wide variety of topics ranging from well established topics to more novel topics, such as large-scale data in cloud computing.

The Panel Chairs, Marianne Winslett, University of Illinois at Urbana-Champaign, and Jeffrey Xu Yu, The Chinese University of Hong Kong, selected two panels addressing exciting questions in our field. This year's edition of the conference had two Awards Committees.

The Best Paper Award Committee consisted of Paolo Atzeni, Università Roma tre, Elisa Bertino, Purdue University, Christian S. Jensen, Aalborg University, Meral Özsoyoglu, Case Western Reserve University, Kian-Lee Tan, National University of Singapore, Yufei Tao, Chinese University of Hong Kong.

The 10-Ten-Year Best Paper Award Committee consisted of Stefano Ceri, Politecnico di Milano, M. Tamer Özsu, University of Waterloo, Canada, and Kyu-Young Whang, KAIST. The 10-year Best Paper award this year is given to two papers: "Don't be lazy, be consistent: Postgres-R, a New Way to Implement Database Replication" by Bettina Kemme and Gustavo Alonso, and "Evolution of the Web and Implications for an Incremental Crawler" by Junghoo Cho and Hector Garcia-Molina.

Assembling the program of VLDB 2010 has been an effort that has involved a large number of members of our community.

We would like to thank first of all the authors who submitted top quality papers addressing many exciting and important research topics. After all, it is their work which makes the conference program. We are grateful to the PC members and external reviewers for the outstanding job with paper review. We appreciated the quality of the reviews, with many suggestions and comments given to the authors, and the active participation to the paper discussions.

We would also like to deeply thank the proceedings editors, Yi Chen, Arizona State University, USA and Y.C. Tay, National University of Singapore, for all their hard work in assembling the proceedings and making sure that all papers were formatted according to the guidelines.

Last but not least, we would like to thank the VLDB General Co-Chairs, Tok Wang Ling, National University of Singapore, and Krithi Ramamritham, Indian Institute of Technology Bombay, for the constant advice and supervision.

Enjoy VLDB 2010!!

Elisa Bertino, Purdue University, USA
Technical Program Chair

Kian Lee Tan, National University of Singapore, Singapore
Core Database Technology Track Chair

Paolo Atzeni, Universita Roma Tre, Italy
Infrastructure for Information Systems Track Chair

Masaru Kitsuregawa, University of Tokyo, Japan
Experiments and Analyses Track Chair

Philip S. Yu, University of Illinois at Chicago, USA
Industrial Applications and Experience Track Chair

Ralf Hartmut Güting, University of Hagen, Germany
Wang-Chien Lee, The Pennsylvania State University, USA
Ee-Peng Lim, Singapore Management University, Singapore
Xiaofeng Meng, Remin University, China
Demonstrations Track Co-Chairs

H.V. Jagadish, University of Michigan , USA
Editor-in-Chief, Proceedings of the VLDB Endowment

CONFERENCE VENUE

GRAND COPTHORNE WATERFRONT HOTEL
392 HAVELOCK ROAD, SINGAPORE 169663

The Grand Copthorne Waterfront Hotel Singapore reflects the way most of us would like to live: amongst resort and exquisite lifestyle decor and a relaxing atmosphere along with great comfort. As one of Asia's premier 5-star choices, the 574-room hotel is nestled next to the Singapore River and close to the Central Business District and Orchard Road. Lifestyle and business-enabling conveniences are offered to facilitate travelers' needs with health and leisure facilities, executive accommodation, wireless connectivity and business support services. Just as your sense of comfort is appeased, your sense of taste in this luxury hotel in Singapore will reach new peaks with dishes from our award-winning restaurants.

CONFERENCE FLOOR PLAN

Layout of Rooms in Level 4 (Grand ballroom for 2 main paper sessions, poster reception, vendor exhibition, registration and teas)

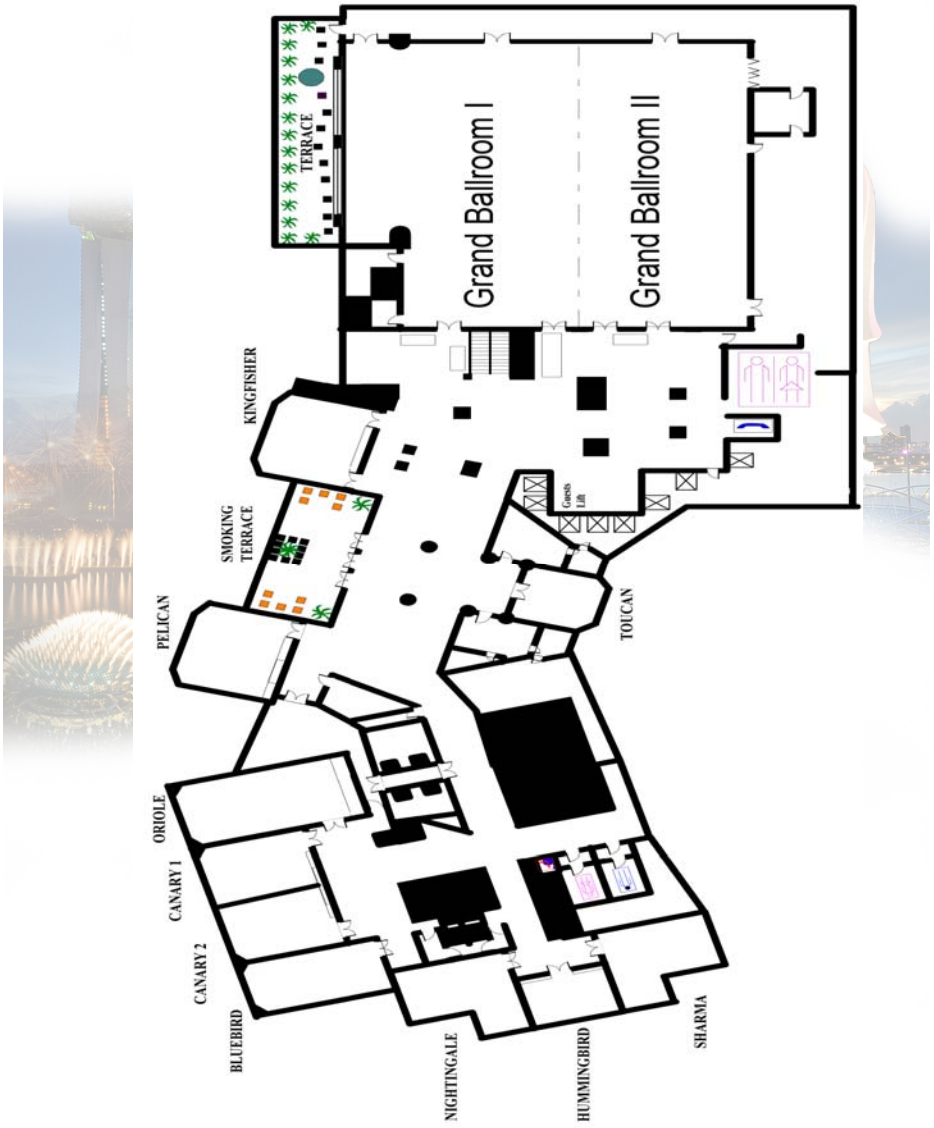

Layout of Rooms in Level 3 (Rooms for 3rd paper session, tutorials, and workshops)

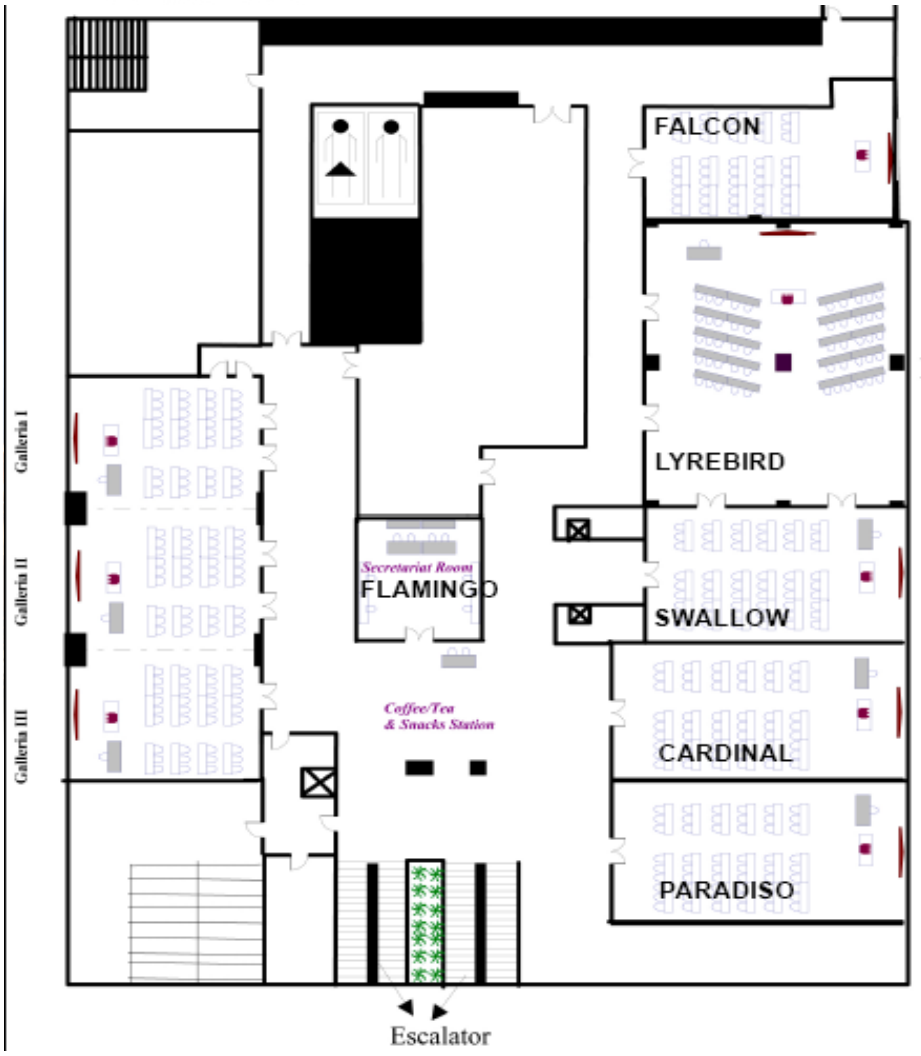

GENERAL INFORMATION

Registration Desk Hours

Sunday 12 September **5.00 pm to 8.00 pm**
Monday 13 September **8.00 am to 8.00 pm**
Tuesday 14 September **8.00 am to 5.00 pm**
Wednesday 15 September **8.00 am to 5.00 pm**
Thursday 16 September **8.00 am to 4.00 pm**
Friday 17 September **8.00 am to 2.00 pm**

Name Tags

Please wear your name tag at all times during the conference, including the welcome reception and conference dinner. You may be asked to present your nametag.

Note to Speakers

If you are scheduled to present, please ensure your PowerPoint is loaded well in advance of your presentation time. A central computer and technician is available at the Registration desk where you may upload your presentation, which in turn will be loaded onto the computer in the appropriate room. Please visit the registration desk if you have any questions or for further details.

Special Dietary Requests

If you have specified vegetarian or halal dietary requirement when registering online, please inform the waiting staff at the conference banquet so that you can be seated at the appropriate table.

Urgent Messages

Urgent messages for delegates can be directed to the registration desk. Messages will be held at the registration desk for collection and the recipient will be notified via a notice board.

Mobile Phones, Pagers & Laptop Sound

As a courtesy to presenters and colleagues, please ensure that all mobile phones, pagers and sound from your laptop are switched off during the conference sessions.

Money

Singapore's unit of currency is the Singapore dollar. Coins have values of 10, 20, and 50 cents, and $1. Notes have a value of $1, $2, $5, $10, $50, and $100. Foreign currency can easily be exchanged at banks, most hotels and money changer booths located at the international airport and throughout Singapore. Major credit cards are accepted throughout Singapore, with travellers cheques accepted in hotels, banks and some stores. Goods and Services Tax (GST) of 7% is charged on all retail goods.

Electricity

Electricity is supplied throughout the country at 230/240 volts, 50 hertz. Most hotels and motels provide 110 volt sockets for use with electric shavers only. For equipment from countries that use 110 volts, an adapter/converter will be necessary unless your equipment is able to run on both 110 and 230/240 volts.

Tipping

Tipping is not obligatory in Singapore, even in restaurants or bars – but it's not frowned upon either. Tipping is left entirely to your discretion, and may be appropriate if you receive excellent service.

SOCIAL EVENTS

VLDB Welcome Reception

Monday (September 13), 19:00
Location: Grand Copthorne Waterfront Hotel, Waterfront Promenade

Conference Banquet

Wednesday (September 15), 18:30
Location: Shangri-la Hotel Ballroom

PROGRAM AT A GLANCE

MONDAY AT A GLANCE							SEPT 13
Room	Galleria I	Galleria II	Galleria III	Cardinal	Falcon	Swallow	Lyrebird
9:00-10:30	Workshop-1 QDB	Workshop-2 MUD	Workshop-3 ADMS	Workshop-4 PersDB	Workshop-5 Ph.D. Workshop	Workshop-6 DMSN	Workshop-7 BIRTE
10:30-11:00	Coffee Break						
Room	Galleria I	Galleria II	Galleria III	Cardinal	Falcon	Swallow	Lyrebird
11:00-12:30	Workshop-1 QDB	Workshop-2 MUD	Workshop-3 ADMS	Workshop-4 PersDB	Workshop-5 Ph.D. Workshop	Workshop-6 DMSN	Workshop-7 BIRTE
12:30-14:00	Lunch ☺ Kiwi Lounge (Level 2)						
Room	Galleria I	Galleria II	Galleria III	Cardinal	Falcon	Swallow	Lyrebird
14:00-15:00	Workshop-1 QDB	Workshop-2 MUD	Workshop-3 ADMS	Workshop-4 PersDB	Workshop-5 Ph.D. Workshop	Workshop-6 DMSN	Workshop-7 BIRTE
15:30-16:00	Coffee Break						
Room	Galleria I	Galleria II	Galleria III	Cardinal	Falcon	Swallow	Lyrebird
16:00-18:00	Workshop-1 QDB	Workshop-2 MUD	Workshop-3 ADMS	Workshop-4 PersDB	Workshop-5 Ph.D. Workshop	Workshop-6 DMSN	Workshop-7 BIRTE
19:00-21:00	VLDB RECEPTION (Waterfront Promenade, Grand Copthorne Waterfront Hotel)						

TUESDAY AT A GLANCE							SEPT 14
8:30-9:00	VLDB Conference Opening (Grand Ballroom)						
9:00-10:30	VLDB Keynote Talk 1: Divesh Srivastava – Enabling Real Time Data Analysis (Grand Ballroom)						
10:30-11:00	Cofffee Break						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
11:00-12:30	Industry-1 Real-Time and Stream Processing	Tutorial-1 Distributed Caching Platform	Research-1 Database Security	Research-2 Parallel and Distributed Databases	Research-3 Data Exchange	Research-4 Database Services and Applications	Demo-1 Core Database Technologies
12:30-14:00	Lunch ☺ Kiwi Lounge (Level 2)						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
14:00-15:30	Research-5 Data Models and Languages	Tutorial-2 Big Data and Cloud Computing: New Wine or Just New Bottles?	Research-6 Semantics	Research-7 Stream Databases	Research-8 RDF and Graphs	Research-9 Middleware Platforms for Data Management	Demo-2 Web Data Management and Information Retrieval Systems
15:30-16:00	Coffee Break						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
16:00-18:00	Research-10 Novel / Advanced Applications	Tutorial-2 Big Data and Cloud Computing: New Wine or Just New Bottles?	Research-11 Ranking Queries	Research-12 Spatial and Temporal Databases	Research-13 Record Linkage	Research-14 Experimental Analysis and Performance	Demo-3 Data Extraction, Integration and Mining

WEDNESDAY AT A GLANCE							SEPT 15
8:45-9:00	Announcement						
9:00-10:30	VLDB Keynote Talk 2: Paul Matsudaira – High-End Biological Imaging Generates Very Large 3D+ and Dynamic Datasets (Grand Ballroom)						
10:30-11:00	Cofffee Break						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
11:00-12:30	Industry-2 Advances in Database Systems	Tutorial-3 Techniques for Similarity Searching in Multimedia Databases	Research-15 Cloud Computing	Research-16 Query Processing I	Research-17 Data Extraction	Research-18 Privacy	Demo-4 Core Database Technologies
12:30-14:00	Lunch ☺ Kiwi Lounge (Level 2)						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
14:00-15:30	Panel-1 Cloud Databases: What's New?	Tutorial-4 Event Processing – Past, Present, Future	Research-19 Probabilistic and Uncertain Databases	Research-20 Databases on Modern Hardware	Research-21 Data Mining	Research-22 Moving Object Databases	Demo-5 Web Data Management and Information Retrieval Systems
15:30-16:00	Coffee Break						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
16:00-17:00	Research-23 Probabilistic Data	Tutorial-4 Event Processing – Past, Present, Future	Research-24 Fuzzy, Probabilistic and Approximate Databases	Research-25 Discovery and Exploration	Research-26 Information Filtering and Dissemination	NO SESSION	Demo-6 Data Extraction, Integration and Mining
17:30-18:00	Leaving the Conference Site for Banquet						
18:30-22:00	VLDB BANQUET (Shangri-la Hotel Ballroom)						

THURSDAY AT A GLANCE							SEPT 16
8:45-10:30	VLDB Awards						
9:00-10:30	Invitted Talks: 10 Year Best Paper Awards Winners (Grand Ballroom)						
10:30-11:00	Cofffee Break						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
11:00-12:30	Industry-3 Data Warehousing and Mining	Tutorial-5 Similarity Search and Mining in Uncertain Databases	Research-27 Query Processing II	Research-28 XML Data	Research-29 Workflows, Transactions and Business Processes	Research-30 Scientific Databases	Research-31 Mobility and Spatial Queries
12:30-14:00	Lunch ☺ Kiwi Lounge (Level 2)						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
14:00-15:30	Industry-4 Mobile and Social Applications	Tutorial-6 Data Management and Mining in Internet Ad Systems	Research-32 Data Anonymization Techniques	Research-33 Querying and Integrating Probabilistic Databases	Research-34 Database Design	Research-35 Query Optimization	Research-36 Graph and Pattern Matching
15:30-16:00	Coffee Break						
Room	Grand Ballroom I	Grand Ballroom II	Galleria Ballroom I	Cardinal	Falcon	Swallow	Lyrebird
16:00-18:00	Panel-2 Time for Our Field to Grow Up	Tutorial-6 Data Management and Mining in Internet Ad Systems	Research-37 Indexing Techniques	Research-38 Query Processing III	Research-39 Streaming and Sensor Data	Research-40 Information Integration and Retrieval	Research-41 Data Mining, Copy Detection, & Data Publishing
18:00	END OF VLDB PROGRAM						

FRIDAY AT A GLANCE				SEPT 17
Room	Cardinal	Falcon	Swallow	Lyrebird
9:00-10:30	Workshop-8 TPCTC	Workshop-9 XSym	Workshop-10 SDM	Workshop-11 SemData
10:30-11:00	Coffee Break			
Room	Cardinal	Falcon	Swallow	Lyrebird
11:00-12:30	Workshop-8 TPCTC	Workshop-9 XSym	Workshop-10 SDM	Workshop-11 SemData
12:30-14:00	Lunch ☺ Kiwi Lounge (Level 2)			
Room	Cardinal	Falcon	Swallow	Lyrebird
14:00-15:00	Workshop-8 TPCTC	Workshop-9 XSym	Workshop-10 SDM	Workshop-11 SemData
15:30-16:00	Coffee Break			
Room	Cardinal	Falcon	Swallow	Lyrebird
16:00-18:00	Workshop-8 TPCTC	Workshop-9 XSym	Workshop-10 SDM	Workshop-11 SemData

Notes:

Notes:

CONFERENCE PROGRAM

MONDAY 19:00 – 21:00

VLDB Reception
Location: Waterfront Promenade, Grand Copthorne Waterfront Hotel

TUESDAY 8:30 – 9:00

VLDB Conference Opening Remarks
Location: Grand Ballroom

TUESDAY 9:00 – 10:30

Keynote 1 – Enabling Real Time Data Analysis
Location: Grand Ballroom
Session Chair: Tok Wang Ling
Speaker: *Divesh Srivastava (AT & T Labs-Research, USA)*

TUESDAY 11:00 – 12:30

Research Session 1: Database Security
Location: Grand Ballroom I
Session Chair: George Kollios

Building Disclosure Risk Aware Query Optimizers for Relational Databases
Mustafa Canim (University of Texas at Dallas, USA), Murat Kantarcioglu (University of Texas at Dallas, USA), Bijit Hore (University of California Irvine, USA), Sharad Mehrotra (University of California Irvine, USA)

Secure Personal Data Servers: a Vision Paper
Tristan Allard (University of Versailles, France), Nicolas Anciaux (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, USA), Indrakshi Ray (Colorado State University, USA), Shaoyi Yin (Institut National de Recherche en Informatique et Automatique, France)

PolicyReplay: Misconfiguration-Response Queries for Data Breach Reporting
Daniel Fabbri (University of Michigan, USA), Kristen LeFevre (University of Michigan, USA), Qiang Zhu (University of Michigan, USA)

TUESDAY 11:00 – 12:30

Research Session 2: Parallel and Distributed Databases
Location: Cardinal
Session Chair: Alan Fekete

Schism: a Workload-Driven Approach to Database Replication and Partitioning
Carlo Curino (Massachusetts Institute of Technology, USA), Yang Zhang (Massachusetts Institute of Technology, USA), Evan Jones (Massachusetts Institute of Technology, USA), Samuel Madden (Massachusetts Institute of Technology, USA)

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)

The Case for Determinism in Database Systems
Alexander Thomson (Yale University, USA), Daniel Abadi (Yale University, USA)

TUESDAY 11:00 – 12:30

Research Session 3: Data Exchange
Location: Falcon
Session Chair: Yannis Velegrakis

MapMerge: Correlating Independent Schema Mappings
Bogdan Alexe (University of California Santa Cruz, USA), Mauricio Hernández (IBM Research, USA), Lucian Popa (IBM Almaden Research Center, USA), Wang-Chiew Tan (University of California Santa Cruz, USA)

Chase Termination: A Constraints Rewriting Approach
Francesca Spezzano (Università della Calabria, Italy), Sergio Greco (Università della Calabria, Italy)

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)

TUESDAY 11:00 – 12:30

Research Session 4: Database Services and Applications
Location: Swallow
Session Chair: Jens Dittrich

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 (ESRI, USA), Yehoshua Sagiv (Hebrew University of Jerusalem, Israel)

Energy Management for MapReduce Clusters

Willis Lang (University of Wisconsin-Madison, USA), Jignesh Patel (University of Wisconsin-Madison, USA)

Toward Scalable Keyword Search over Relational Data

Akanksha Baid (University of Wisconsin-Madison, USA), Ian Rae (University of Wisconsin-Madison, USA), Jiexing Li (University of Wisconsin-Madison, USA), AnHai Doan (University of Wisconsin-Madison, USA), Jeffrey Naughton (University of Wisconsin-Madison, USA)

TUESDAY 11:00 – 12:30

Industry Session 1: Real-Time and Stream Processing

Location: Grand Ballroom-1

Session Chair: Hui-I Hsiao

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), Michalis Vlachos (IBM Research - Zurich, Switzerland)

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, USA), Martin Hirzel (IBM Thomas J. Watson Research Center, USA), Henrique Andrade (IBM Thomas J. Watson Research Center, USA), Buğra Gedik (IBM Thomas J. Watson Research Center, USA), Kun-Lung Wu (IBM Thomas J. Watson Research Center, USA)

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 Chai (National University of Singapore, Republic of Singapore)

TUESDAY 11:00 – 12:30

Tutorial 1: Distributed Caching Platforms

Location: Grand Ballroom-II

Presenter: Anil Nori (Microsoft)

TUESDAY 11:00 – 12:30

Demo Session 1: Core Database Technologies

Location: Lyrebird

The Picasso Database Query Optimizer Visualizer

Jayant Haritsa (Indian Institute of Science, India)

CODS: Evolving Data Efficiently and Scalably in Column Oriented Databases

Ziyang Liu (Arizona State University, USA), Sivaramakrishnan Natarajan (Arizona State University, USA), Bin He (IBM Almaden Research Center, USA), Hui-I Hsiao (IBM Almaden Research Center, USA), Yi Chen (Arizona State University, USA)

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)

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

Justin Levandoski (University of Minnesota, USA), Mohamed Mokbel (University of Minnesota, USA), Mohamed Khalefa (University of Minnesota, USA)

Cloudy: A Modular Cloud Storage System

Donald Kossmann (Eidgenössische Technische Hochschule Zürich, Switzerland), Tim Kraska (University of California Berkeley, USA), 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)

Geospatial Stream Query Processing using Microsoft SQL Server StreamInsight

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

Using XMorph to Transform XML Data

Curtis Dyreson (Utah State University, USA), Sourav S Bhowmick (Nanyang Technological University, Republic of Singapore), Kirankanth Mallampalli (Utah State University, USA)

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

Di Wang (Worcester Polytechnic Institute, USA), Elke Rundensteiner (Worcester Polytechnic Institute, USA), Richard Ellison (University of Massachusetts Medical School, USA), Han Wang (Worcester Polytechnic Institute, USA)

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)

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)

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, USA), Fan Ping (State University of New York University at Albany, USA), Jeong-Hyon Hwang (State University of New York University at Albany, USA)

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)

TUESDAY 14:00 – 15:30

Research Session 5: Data Models and Languages
Location: Grand Ballroom I
Session Chair: Stefano Paraboschi

From Regular Expressions to Nested Words: Unifying Languages and Query Execution for Relational and XML Sequences
Barzan Mozafari (University of California Los Angeles, USA), Kai Zeng (University of California Los Angeles, USA), Carlo Zaniolo (University of California Los Angeles, USA)

Avalanche-Safe LINQ Compilation
Torsten Grust (Universität Tübingen, Germany), Jan Rittinger (Universität Tübingen, Germany), Tom Schreiber (Universität Tübingen, Germany)

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), Wenyuan Yu (University of Edinburgh, United Kingdom)

TUESDAY 14:00 – 15:30

Research Session 6: Semantics
Location: Galleria Ballroom I
Session Chair: Divesh Srivastava

Explaining Missing Answers to SPJUA Queries
Melanie Herschel (Universität Tübingen, Germany), Mauricio Hernández (IBM Research, USA)

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, USA)

Evaluating Entity Resolution Results
David Menestrina (Stanford University, USA), Steven Whang (Stanford University, USA), Hector Garcia-Molina (Stanford University, USA)

TUESDAY 14:00 – 15:30

Research Session 7: Stream Databases
Location: Cardinal
Session Chair: Rajeev Rastogi

High-Performance Dynamic Pattern Matching over Disordered Streams
Badrish Chandramouli (Microsoft Research, USA), Jonathan Goldstein (Microsoft Research, USA), David Maier (Portland State University, USA)

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, USA), Renée Miller

(University of Toronto, Canada), Nesime Tatbul (Eidgenössische Technische Hochschule Zürich, Switzerland)

Recognizing Patterns in Streams with Imprecise Timestamps

Haopeng Zhang (University of Massachusetts Amherst, USA), Yanlei Diao (University of Massachusetts Amherst, USA), Neil Immerman (University of Massachusetts Amherst, USA)

TUESDAY 14:00 – 15:30

Research Session 8: RDF and Graphs

Location: Falcon

Session Chair: Wook-Shin Han

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)

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)

GRAIL: Scalable Reachability Index for Large Graphs

Hilmi Yildirim (Rensselaer Polytechnic Institute, USA), Vineet Chaoji (Yahoo! Research Labs, USA), Mohammed Zaki (Rensselaer Polytechnic Institute, USA)

TUESDAY 14:00 – 15:30

Research Session 9: Middleware Platforms for Data Management

Location: Swallow

Session Chair: Mike Carey

HaLoop: Efficient Iterative Data Processing on Large Clusters

Yingyi Bu (University of Washington, USA), Bill Howe (University of Washington, USA), Magdalena Balazinska (University of Washington, USA), Michael Ernst (University of Washington, USA)

The Impact of Virtual Views on Containment

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

Updatable and Evolvable Transforms for Virtual Databases

James Terwilliger (Microsoft Corporation, USA), Lois Delcambre (Portland State University, USA), David Maier (Portland State University, USA), Jeremy Steinhauer (Portland State University, USA), Scott Britell (Portland State University, USA)

TUESDAY 14:00 – 15:30

Tutorial 2: Big Data and Cloud Computing: New Wine or Just New Bottles?
Location: Grand Ballroom-II

Presenters: Divyakant Agrawal (University of California, Santa Barbara, USA), Sudipto Das (University of California, Santa Barbara, USA), Amr El Abbadi (University of California, Santa Barbara, USA)

TUESDAY 14:00 – 15:30

Demo Session 2: Web Data Management and Information Retrieval Systems
Location: Lyrebird

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, USA), Lizhu Zhou (Tsinghua University, People's Republic of China)

TimeTrails: A System for Exploring Spatio-Temporal Information in Documents
Jannik Strötgen (University of Heidelberg, Germany), Michael Gertz (University of Heidelberg, Germany)

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)

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)

XSACT: A Comparison Tool for Structured Search Results
Ziyang Liu (Arizona State University, USA), Sivaramakrishnan Natarajan (Arizona State University, USA), Peng Sun (Arizona State University, USA), Stephen Booher (Arizona State University, USA), Tim Moohan (Arizona State University, USA), Robert Winkler (Arizona State University, USA), Yi Chen (Arizona State University, USA)

ObjectRunner: Lightweight, Targeted Extraction and Querying of Structured Web Data
Talel Abdessalem (Télécom Telecom ParisTech, France), Bogdan Cautis (Télécom ParisTech, France), Nora Derouiche (Télécom ParisTech, France)

ROXXI: Reviving witness dOcuments to eXplore 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)

EXTRUCT: Using Deep Structural Information in XML Keyword Search
Arash Termehchy (University of Illinois at Urbana-Champaign, USA), Marianne Winslett (University of Illinois at Urbana-Champaign, USA)

SQL QueRIE Recommendations
Javad Akbarnejad (San Jose State University, USA), Gloria Chatzopoulou (University of California Riverside, USA), Magdalini Eirinaki (San Jose State University, USA), Suju Koshy (San Jose State University, USA), Sarika Mittal (San Jose State University, USA), Duc On (San Jose State University, USA), Neoklis Polyzotis (University of California Santa Cruz, USA), Jothi Swarubini Vindhiya Varman (San Jose State University, USA)

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)

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)

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

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

TUESDAY 16:00 – 17:30

Research Session 10: Novel/Advanced Applications

Location: Grand Ballroom

Session Chair: Johann-Christoph Freytag

Navigating in Complex Mashed-Up Applications

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

Dremel: Interactive Analysis of Web-Scale Datasets

Sergey Melnik (Google, USA), Andrey Gubarev (Google, USA), Jing Jing Long (Google, USA), Geoffrey Romer (Google, USA), Shiva Shivakumar (Google, USA), Matt Tolton (Google, USA), Theo Vassilakis (Google, USA)

On Graph Query Optimization in Large Networks

Peixiang Zhao (University of Illinois at Urbana-Champaign, USA), Jiawei Han (University of Illinois at Urbana-Champaign, USA)

TUESDAY 16:00 – 17:30

Research Session 11: Ranking Queries

Location: Galleria Ballroom I

Session Chair: Alfons Kemper

Proximity Rank Join

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

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

Akrivi Vlachou (Norwegian University of Science and Technology, Norway), Christos Doulkeridis (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)

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)

TUESDAY 16:00 – 17:30

Research Session 12: Spatial and Temporal Databases
Location: Cardinal
***Session Chair*: Cyrus Shehabi**

Parsimonious Linear Fingerprinting for Time Series
Lei Li (Carnegie Mellon University, USA), B. Aditya Prakash (Carnegie Mellon University, USA), Christos Faloutsos (Carnegie Mellon University, USA)

The HV-tree: a Memory Hierarchy Aware Version Index
Rui Zhang (University of Melbourne, Australia), Martin Stradling (University of Melbourne, Australia)

Transforming Range Queries To Equivalent Box Queries To Optimize Page Access
Sakti Pramanik (Michigan State University, USA), Alok Watve (Michigan State University, USA), Chad Meiners (Michigan State University, USA), Alex Liu (Michigan State University, USA)

TUESDAY 16:00 – 17:30

Research Session 13: Record Linkage
Location: Falcon
***Session Chair*: Elisa Bertino**

Record Linkage with Uniqueness Constraints and Erroneous Values
Songtao Guo (AT&T Interactive, USA), Xin Dong (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs -Research, USA), Remi Zajac (AT&T Interactive, USA)

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)

Behavior Based Record Linkage
Mohamed Yakout (Purdue University, USA), Ahmed Elmagarmid (Purdue University, USA), Hazem Elmeleegy (Purdue University, USA), Mourad Ouzzani (Purdue University, USA), Alan Qi (Purdue University, USA)

TUESDAY 16:00 – 18:00

Research Session 14: Experimental Analysis and Performance
Location: Swallow
***Session Chair*: Paul Larson**

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)

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)

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)

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)

TUESDAY 16:00 – 17:30

Tutorial 2: Big Data and Cloud Computing: New Wine or Just New Bottles?
Location: Grand Ballroom-II

Presenters: Divyakant Agrawal (University of California, Santa Barbara, USA), Sudipto Das (University of California, Santa Barbara, USA), Amr El Abbadi (University of California, Santa Barbara, USA)

TUESDAY 16:00 – 17:30

Demo Session 3: Data Extraction, Integration and Mining
Location: Lyrebird

Deep Web Integration with VisQI
Thomas Kabisch (Humboldt Universität zu Berlin, Germany), Eduard Dragut (University of Illinois at Chicago, USA), Clement Yu (University of Illinois at Chicago, USA), Ulf Leser, Humboldt Universität zu Berlin, Germany)

SOLOMON: Seeking the Truth Via Copying Detection
Xin Dong (AT&T Labs - Research, USA), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

Just-in-time Data Integration in Action
Martin Hentschel (Eidgenössische Technische Hochschule Zürich, Switzerland), Laura Haas (IBM Almaden Research Center, USA), Renée Miller (University of Toronto, Canada)

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 Heimel (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)

Using Sentinel Technology in the TARGIT BI Suite
Morten Middelfart (TARGIT A/S, Denmark), Torben Bach Pedersen (Aalborg University, Denmark)

CoDA: Interactive Cluster Based Concept Discovery
Stephan Günnemann (RWTH Aachen University, Germany), Ines Färber (RWTH Aachen University, Germany), Hardy Kremer (RWTH Aachen University, Germany), Thomas Seidl (RWTH Aachen University, Germany)

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 Velegrakis (University of Trento, Italy)

Data Auditor: Exploring Data Quality and Semantics using Pattern Tableaux

Lukasz Golab (AT&T Labs - Research, USA), Howard Karloff (AT&T Labs - Research, USA), Flip Korn (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

WEDNESDAY 8:45 – 9:00

Announcement
Location: Grand Ballroom

WEDNESDAY 9:00 – 10:30

Keynote 2 – High-End Biological Imaging Generates Very Large 3D+ and Dynamic Datasets
Location: Grand Ballroom
Session Chair: Kian-Lee Tan
Speaker: *Paul Matsudaira (National University of Singapore, Singapore)*

WEDNESDAY 11:00 – 12:30

Research Session 15: Cloud Computing
Location: Galleria Ballroom I
Session Chair: Karl Aberer

MRShare: Sharing Across Multiple Queries in MapReduce
Tomasz Nykiel (University of Toronto, Canada), Michalis Potamias (Boston University, USA), Chaitanya Mishra (Facebook, USA), George Kollios (Boston University, USA), Nick Koudas (University of Toronto, Canada)

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)

Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)
Jens Dittrich (Saarland University, Germany), Jorge-Arnulfo Quiane-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 (Max-Planck Institute for Informatics, Germany), Jörg Schad (Saarland University, Germany)

WEDNESDAY 11:00 – 12:30

Research Session 16: Query Processing I
Location: Cardinal
Session Chair: S Sudarshan

Slicing Long-Running Queries
Nicolas Bruno (Microsoft Research, USA), Vivek Narasayya (Microsoft Research, USA), Ravi Ramamurthy (Microsoft Research, USA)

Sharing-Aware Horizontal Partitioning for Exploiting Correlations during Query Processing
Kostas Tzoumas (Aalborg University, Denmark), Amol Deshpande (University of Maryland College Park, USA), Christian Jensen (Aarhus University, Denmark)

Advanced Processing for Ontological Queries

Andrea Calì (University of Oxford, United Kingdom), Georg Gottlob (University of Oxford, United Kingdom), Andreas Pieris (Brunei University, United Kingdom)

WEDNESDAY 11:00 – 12:30

Research Session 17: Data Extraction

Location: Falcon

Session Chair: Jiaheng Lu

Towards The Web of Concepts: Extracting Concepts from Large Datasets

Aditya Parameswaran (Stanford University, USA), Hector Garcia-Molina (Stanford University, USA), Anand Rajaraman (Kosmix Corporation, USA)

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)

Automatic Rule Refinement for Information Extraction

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

WEDNESDAY 11:00 – 12:30

Research Session 18: Privacy

Location: Swallow

Session Chair: Panagiotis Karras

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)

Small Domain Randomization: Same Privacy, More Utility

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

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, USA), Dimitris Papadias (The Hong Kong University of Science and Technology, People's Republic of China)

WEDNESDAY 11:00 – 12:30

Industry Session 2: Advances in Database Systems
Location: Grand Ballroom-I
Session Chair: Ilyas Ehab

FlashStore: High Throughput Persistent Key-Value Store
Biplob Debnath (University of Minnesota, USA), Sudipta Sengupta (Microsoft Research, Redmond, USA), Jin Li (Microsoft Research, Redmond, USA)

MEET DB2: Automated Database Migration Evaluation
Reynold Xin (University of California Berkeley, USA), Patrick Dantressangle (IBM Hursley, United Kingdom), Sam Lightstone (IBM Research, Canada), William McLaren (IBM, USA), Steve Schormann (IBM Research, Canada), Maria Schwenger (IBM, USA)

SSD Bufferpool Extensions for Database Systems
Mustafa Canim (University of Texas at Dallas, USA), George Mihaila (IBM Thomas J. Watson Research Center, USA), Bishwaranjan Bhattacharjee (IBM Thomas J. Watson Research Center, USA), Kenneth Ross (Columbia University, USA), Christian Lang (IBM Thomas J. Watson Research Center, USA)

WEDNESDAY 11:00 – 12:30

Tutorial 3: Techniques for Similarity Searching in Multimedia Databases
Location: Grand Ballroom-II

Presenter: Hanan Samet (University of Maryland College Park, USA)

WEDNESDAY 11:00 – 12:30

Demo Session 4: Core Database Technologies
Location: Lyrebird

The program is same as that of Demo Session 1 (Tuesday, 11:00 – 12:30)

WEDNESDAY 14:00 – 15:30

Research Session 19: Probabilistic and Uncertain Databases
Location: Galleria Ballroom I
Session Chair: Reynold Cheng

UPI: A Primary Index for Uncertain Databases
Hideaki Kimura (Brown University, USA), Samuel Madden (Massachusetts Institute of Technology, USA), Stan Zdonik (Brown University, USA)

Ranking Continuous Probabilistic Datasets
Jian Li (University of Maryland College Park, USA), Amol Deshpande (University of Maryland College Park, USA)

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)

WEDNESDAY 14:00 – 15:30

Research Session 20: Databases on Modern Hardware

Location: Cardinal

Session Chair: Anthony Tung

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)

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)

Aether: A Scalable Approach to Logging

Ryan Johnson (Carnegie Mellon University, USA), Ippokratis Pandis (Carnegie Mellon University, USA), Radu Stoica (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Manos Athanassoulis (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland)

WEDNESDAY 14:00 – 15:30

Research Session 21: Data Mining

Location: Falcon

Session Chair: Maurice van Keulen

Scalable Discovery of Best Clusters on Large Graphs

Kathy Macropol (University of California Santa Barbara, USA), Ambuj Singh (University of California Santa Barbara, USA)

An Architecture for Parallel Topic Models

Alexander Smola (Yahoo! Research Labs, USA), Shravan Narayanamurthy (Yahoo! Research Labs, India)

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

Dong Xin (Microsoft Research, USA), Yeye He (University of Wisconsin-Madison, USA), Venkatesh Ganti (Google, USA)

WEDNESDAY 14:00 – 15:30

Research Session 22: Moving Object Databases
Location: Swallow
Session Chair: Christian S Jensen

Swarm: Mining Relaxed Temporal Moving Object Clusters
Zhenhui Li (University of Illinois at Urbana-Champaign, USA), Bolin Ding (University of Illinois at Urbana-Champaign, USA), Jiawei Han (University of Illinois at Urbana-Champaign, USA), Roland Kays (New York State Museum, USA)

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)

Shortest Path Computation on Air Indexes
Georgios Kellaris (Singapore Management University, Republic of Singapore), Kyriakos Mouratidis (Singapore Management University, Republic of Singapore)

WEDNESDAY 14:00 – 15:30

Tutorial 4: Event Processing: Past, Present, Future?
Location: Grand Ballroom II

Presenter: Opher Etzion (IBM Haifa Research Lab, Israel)

WEDNESDAY 14:00 – 15:30

Demo Session 5: Web Data Management and Information Retrieval Systems
Location: Lyrebird

The program is same as that of Demo Session 2 (Tuesday, 14:00 – 15:30)

WEDNESDAY 14:00 – 15:30

Panel 1: Cloud Databases: What's New?
Location: Grand Ballroom I

Moderator: Raghu Ramakrishnan (Yahoo! Research)
Members: Daniel Abadi (Yale University, USA), Michael Carey (University of California, Irwine, USA), Surajit Chaudhuri (Microsoft Research, USA), Hector Garcia-Molina (Stanford University, USA), Jignesh M Patel (University of Wisconsin, Madison, USA)

WEDNESDAY 16:00 – 17:00

Research Session 23: Probabilistic Data
Location: Grand Ballroom I
Session Chair: Mauricio Hernandez

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)

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)

WEDNESDAY 16:00 – 17:00

Research Session 24: Fuzzy, Probabilistic and Approximate Databases
Location: Galleria Ballroom I
Session Chair: Lei Chen

MCDB-R: Risk Analysis in the Database
Peter Haas (IBM Research, USA), Christopher Jermaine (Rice University, USA), Subi Arumugam (University of Florida Gainesville, USA), Fei Xu (Microsoft Corporation, USA), Luis Perez (Rice University, USA), Ravi Jampani (University of Florida Gainesville, USA)

Scalable Probabilistic Databases with Factor Graphs and MCMC
Michael Wick (University of Massachusetts Amherst, USA), Andrew McCallum (University of Massachusetts Amherst, USA), Gerome Miklau (University of Massachusetts Amherst, USA)

WEDNESDAY 16:00 – 17:00

Research Session 25: Discovery and Exploration
Location: Cardinal
Session Chair: Erhard Rahm

On Multi-Column Foreign Key Discovery
Meihui Zhang (National University of Singapore, Republic of Singapore), Marios Hadjieleftheriou (AT&T Labs - Research, USA), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Cecilia Procopiuc (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

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)

WEDNESDAY 16:00 – 17:00

Research Session 26: Information Filtering and Dissemination
Location: Falcon
Session Chair: Themis Palpanas

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)

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

WEDNESDAY 16:00 – 17:30

Tutorial 4: Event Processing: Past, Present, Future?
Location: Grand Ballroom II

Presenter: Opher Etzion (IBM Haifa Research Lab, Israel)

WEDNESDAY 16:00 – 17:30

Demo Session 6: Data Extraction, Integration and Mining
Location: Lyrebird

The program is same as that of Demo Session 3 (Tuesday, 16:00 – 17:30)

WEDNESDAY 18:30 – 22:00

Banquet (to leave the conference site between 17:00 to 18:00)
Location: Shangri-la Hotel Ballroom

THURSDAY 8:45 – 9:00

VLDB Awards Ceremony
Location: Grand Ballroom

THURSDAY 9:00 – 9:45

10 Year Best Paper Award Keynote 1: Dealing With Web Data: History and Look Ahead
Location: Grand Ballroom
Session Chair: Paolo Atzeni

Presenters:
Junghoo Choo (University of California, Los Angeles, USA), Hector Garcia-Molina (Stanford University, USA)

THURSDAY 9:45 – 10:30

10 Year Best Pape Award Keynote 2: Database Replication: A Tale of Research Across Communities
Location: Grand Ballroom
Session Chair: Masaru Kitsuregawa

Presenters: Bettima Kemme (McGill University, Canada), Gustavo Alonso (ETH-Zurich, Switzerland)

THURSDAY 11:00 – 12:30

Research Session 27: Query Processing II
Location: Galleria Ballroom I
Session Chair: Jayant Haritsa

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)

Processing Top-k Join Queries
Minji Wu (Rutgers University, USA), Laure Berti-Equille (University of Rennes 1, France), Amélie Marian (Rutgers University, USA), Cecilia Procopiuc (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

Two-way Replacement Selection
Xavier Martinez-Palau (Universitat Politècnica de Catalunya, Spain), David Dominguez-Sal (Universitat Politècnica de Catalunya, Spain), Josep Lluis Larriba-Pey (Universitat Politècnica de Catalunya, Spain)

THURSDAY 11:00 – 12:30

Research Session 28: XML Data
Location: Cardinal
Session Chair: Fischer Peter Michael

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)

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)

Destabilizers and Independence of XML Updates
Michael Benedikt (University of Oxford, United Kingdom), James Cheney (University of Edinburgh, United Kingdom)

THURSDAY 11:00 – 12:30

Research Session 29: Workflows, Transactions, and Business Processes
Location: Falcon
Session Chair: Z Meral Ozsoyoglu

Searching Workflows with Hierarchical Views
Ziyang Liu (Arizona State University, USA), Qihong Shao (Arizona State University, USA), Yi Chen (Arizona State University, USA)

Data-Oriented Transaction Execution
Ippokratis Pandis (Carnegie Mellon University, USA), Ryan Johnson (Carnegie Mellon University, USA), Nikos Hardavellas (Northwestern University, USA), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland)

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, USA), Tom Yam (Tel Aviv University, Israel)

THURSDAY 11:00 – 12:30

Research Session 30: Scientific Databases
Location: Swallow
Session Chair: Jan Paredaens

Behavioral Simulations in MapReduce
Guozhang Wang (Cornell University, USA), Marcos Vaz Salles (Cornell University, USA), Benjamin Sowell (Cornell University, USA), Xun Wang (Cornell University, USA), Tuan Cao (Cornell University, USA), Alan Demers (Cornell University, USA), Johannes Gehrke (Cornell University, USA), Walker White (Cornell University, USA)

A*-tree: A Structure for Storage and Modeling of Uncertain Multidimensional Arrays
Tingjian Ge (University of Kentucky, USA), Stan Zdonik (Brown University, USA)

On Dense Pattern Mining in Graph Streams
Charu Aggarwal (IBM Thomas J. Watson Research Center, USA), Yao Li (University of Illinois at Chicago, USA), Philip Yu (University of Illinois at Chicago, USA), Ruoming Jin (Kent State University, USA)

THURSDAY 11:00 – 12:30

Research Session 31: Mobility and Spatial Queries
Location: Lyrebird
Session Chair: Kyriakos Mouratidis

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 Saltenis (Aalborg University, Denmark), Kostas Tzoumas (Aalborg University, Denmark)

k-Nearest Neighbors in Uncertain Graphs
Michalis Potamias (Boston University, USA), Francesco Bonchi (Yahoo! Research Labs, USA), Aristides Gionis (Yahoo! Research Labs, USA), George Kollios (Boston University, USA)

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)

THURSDAY 11:00 – 12:30

Industry Session 3: Data Warehousing and Mining
Location: Grand Ballroom I
Session Chair: Elisa Bertino

DataGarage: Warehousing Massive Amounts of Performance Data on Commodity Servers
Charles Loboz (Microsoft Corporation, USA), Slawek Smyl (Microsoft Corporation, USA), Suman Nath (Microsoft Research, USA)

Cheetah: A High Performance, Custom Data Warehouse on Top of MapReduce
Songting Chen (Turn Incorporated, USA)

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 (Ohio State University, USA), Wagner Meira (Universidade Federal de Minas Gerais, Brazil), Srinivasan Parthasarathy (Ohio State University, USA)

THURSDAY 11:00 – 12:30

Tutorial 5: Similarity Search and Mining in Uncertain Databases
Location: Grand Ballroom II

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

THURSDAY 14:00 – 15:30

Research Session 32: Data Anonymization Techniques
Location: Galleria Ballroom I
Session Chair: Murat Kantarcioglu

Boosting the Accuracy of Differentially-Private Histograms Through Consistency
Michael Hay (University of Massachusetts Amherst, USA), Vibhor Rastogi (University of Washington, USA), Gerome Miklau (University of Massachusetts Amherst, USA), Dan Suciu (University of Washington, USA)

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 (National University of Singapore, Republic of Singapore), Kian-Lee Tan (National University of Singapore, Republic of Singapore)

Minimizing Minimality and Maximizing Utility: Analyzing Method-based attacks on Anonymized Data
Graham Cormode (AT&T Labs - Research, USA), Ninghui Li (Purdue University, USA), Tiancheng Li (Purdue University, USA), Divesh Srivastava (AT&T Labs - Research, USA)

THURSDAY 14:00 – 15:30

Research Session 33: Querying and Integrating Probabilistic Databases
Location: Cardinal
Session Chair: Zhenjie Zhang

Querying Probabilistic Information Extraction
Daisy Zhe Wang (University of California Berkeley, USA), Michael Franklin (University of California Berkeley, USA), Minos Garofalakis (Technical University of Crete, Greece), Joseph Hellerstein (University of California Berkeley, USA)

Read-Once Functions and Query Evaluation in Probabilistic Databases
Prithviraj Sen (Yahoo! Research Labs, USA), Amol Deshpande (University of Maryland College Park, USA), Lise Getoor (University of Maryland College Park, USA)

Foundations of Uncertain-Data Integration
Parag Agrawal (Stanford University, USA), Anish Das Sarma (Yahoo! Research Labs, USA), Jeffrey Ullman (Stanford University, USA), Jennifer Widom (Stanford University, USA)

THURSDAY 14:00 – 15:30

Research Session 34: Database Design
Location: Falcon
Session Chair: Lois Delcambre

Identifying, Attributing and Describing Spatial Bursts
Michael Mathioudakis (University of Toronto, Canada), Nilesh Bansal (University of Toronto, Canada), Nick Koudas (University of Toronto, Canada)

CORADD: Correlation Aware Database Designer for Materialized Views and Indexes
Hideaki Kimura (Brown University, USA), George Huo (Google, USA), Alexander Rasin (Brown University, USA), Samuel Madden (Massachusetts Institute of Technology, USA), Stanley Zdonik (Brown University, USA)

Regret-Minimizing Representative Databases
Danupon Nanongkai (Georgia Institute of Technology, USA) , Atish Das Sarma (Georgia Institute of Technology, USA), Ashwin Lall (Georgia Institute of Technology, USA), Richard Lipton (Georgia Institute of Technology, USA), Jun Xu (Georgia Institute of Technology, USA)

THURSDAY 14:00 – 15:30

Research Session 35: Query Optimization
Location: Swallow
Session Chair: Torsten Grust

An Access CostAware Approach for Object Retrieval over Multiple Sources
Benjamin Arai (University of California Riverside, USA), Gautam Das (University of Texas at Arlington, USA), Dimitrios Gunopulos (University of Athens, Greece), Vagelis Hristidis (Florida International University, USA), Nick Koudas (University of Toronto, Canada)

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)

Xplus: A SQL-Tuning-Aware Query Optimizer
Herodotos Herodotou (Duke University, USA), Shivnath Babu (Duke University, USA)

THURSDAY 14:00 – 15:30

Research Session 36: Graph and Pattern Matching
Location: Lyrebird
Session Chair: Jeffrey Xu Yu

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)

SigMatch: Fast and Scalable Multi-Pattern Matching
Ramakrishnan Kandhan (University of Wisconsin-Madison, USA), Nikhil Teletia (University of Wisconsin-Madison, USA), Jignesh Patel (University of Wisconsin-Madison, USA)

SAPPER: Subgraph Indexing and Approximate Matching in Large Graphs
Shijie Zhang (Case Western Reserve University, USA), Jiong Yang (Case Western Reserve University, USA), Wei Jin (Case Western Reserve University, USA)

THURSDAY 14:00 – 15:30

Industry Session 4: Mobile and Social Applications
Location: Grand Ballroom I
Session Chair: Wang-Chien Lee

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)

RoadTrack: Scaling Location Updates for Mobile Clients on Road Networks with Query Awareness
Peter Pesti (Georgia Institute of Technology, USA), Ling Liu (Georgia Institute of Technology, USA), Bhuvan Bamba (Georgia Institute of Technology, USA), Arun Iyengar (IBM Thomas J. Watson Research Center, USA), Matt Weber (Georgia Institute of Technology, USA)

Confucius and its Intelligent Disciples: Integrating Social with Search
Xiance Si (Google Research, People's Republic of China), Edward Y. Chang (Google Research, People's Republic of China), Zoltán Gyöngyi (Google Research, USA), Maosong Sun (Tsinghua University, People's Republic of China)

THURSDAY 14:00 – 15:30

Tutorial 6: Data Management and Mining in Internet Ad Systems
Location: Grand Ballroom II

Presenter: S Muthukrishnan (Rutgers University, USA)

THURSDAY 16:00 – 17:30

Research Session 37: Indexing Techniques
Location: Galleria Ballroom I
Session Chair: Nick Koudas

Tree Indexing on Solid State Drives
Yinan Li (University of Wisconsin-Madison, USA), 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)

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, USA)

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)

THURSDAY 16:00 – 17:30

Research Session 38: Query Processing III
Location: Cardinal
Session Chair: Man Lung Yiu

VoR-Tree: R-trees with Voronoi Diagrams for Efficient Processing of Spatial Nearest Neighbor Queries
Mehdi Sharifzadeh (Google Santa Monica, USA), Cyrus Shahabi (University of Southern California, USA)

Efficient RkNN Retrieval with Arbitrary Non-Metric Similarity Measures
Deepak P (IBM India Research Lab, India), Prasad Deshpande (IBM India Research Lab, India)

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 W. Cheung (The University of Hong Kong, People's Republic of China)

THURSDAY 16:00 – 18:00

Research Session 39: Streaming and Sensor Data
Location: Falcon
Session Chair: Rui Zhang

Achieving High Output Quality under Limited Resources through Structure-based Spilling in XML Streams
Mingzhu Wei (Worcester Polytechnic Institute, USA), Elke Rundensteiner (Worcester Polytechnic Institute, USA), Murali Mani (University of Michigan at Flint, USA)

Dynamic Join Optimization in Multi-Hop Wireless Sensor Networks
Svilen Mihaylov (University of Pennsylvania, USA), Marie Jacob (University of Pennsylvania, USA), Zachary Ives (University of Pennsylvania, USA), Sudipto Guha (University of Pennsylvania, USA)

Database-support for Continuous Prediction Queries over Streaming Data
Mert Akdere (Brown University, USA), Ugur Cetintemel (Brown University, USA), Eli Upfal (Brown University, USA)

Conditioning and Aggregating Uncertain Data Streams: Going Beyond Expectations
Thanh Tran (University of Massachusetts Amherst, USA), Andrew McGregor (University of Massachusetts Amherst, USA), Yanlei Diao (University of Massachusetts Amherst, USA), Liping Peng (University of Massachusetts Amherst, USA), Anna Liu (University of Massachusetts Amherst, USA)

THURSDAY 16:00 – 17:30

Research Session 40: Information Integration and Retrieval
Location: Swallow
Session Chair: Matthias Renz

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 Almaden Research Center, USA)

Entity Resolution with Evolving Rules
Steven Whang (Stanford University, USA), Hector Garcia-Molina (Stanford University, USA)

Annotating and Searching Web Tables Using Entities, Types and Relationships
Girija Limaye (IIT Bombay, India), Sunita Sarawagi (IIT Bombay, India), Soumen Chakrabarti (IIT Bombay, India)

THURSDAY 16:00 – 17:30

Research Session 41: Data Mining, Copy Detection and Data Publication
Location: Lyrebird
Session Chair: Kyuseok Shim

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)

Global Detection of Complex Copying Relationships Between Sources
Xin Dong (AT&T Labs - Research, USA), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

Fragments and Loose Associations: Respecting Privacy in Data Publishing
Sabrina De Capitani di Vimercati (Universita' degli Studi di Milano, Italy), Sara Foresti (Universita' degli Studi di Milano, Italy), Sushil Jajodia (George Mason University Fairfax, United States of America), Stefano Paraboschi (University of Bergamo, Italy), Pierangela Samarati (Universita' degli Studi di Milano, Italy)

THURSDAY 16:00 – 17:30

Tutorial 6: Data Management and Mining in Internet Ad Systems
Location: Grand Ballroom II

Presenter: S Muthukrishnan (Rutgers University, USA)

THURSDAY 16:00 – 17:00

Panel 2: Time for Our Field to Grow Up
Location: Grand Ballroom I

Moderator: Marianne Winslett (University of Illinois, Urbana-Champaign, USA)
Members: Anastassia Ailamaki (EPFL, Switzerland), Laura Haas (IBM Almaden Research Center, USA), H V Jagadish (University of Michigan, USA), David Maier (Portland State University, USA), Tamer Ozsu (University of Waterloo, Canada)

Notes:

Notes:

INVITED TALKS AND KEYNOTES

Keynote Talk 1: Enabling Real Time Data Analysis
Divesh Srivastava, AT&T Lab-Research, USA
14th September (Tuesday), 9:00 AM – 10:30 AM
Location: Grand Ballroom

Network-based services are becoming a ubiquitous part of our lives, to the point where individuals and businesses have often come to critically rely on them. Building and maintaining such reliable, high performance network and service infrastructures requires the ability to rapidly investigate and resolve complex service and performance impacting issues. To achieve this, it is important to collect, correlate and analyze massive amounts of data, such as network element fault logs, configuration and topology information, performance data and traffic measurements, from a diverse collection of data sources in real time.

We have designed and implemented a variety of data systems at AT&T Labs Research to build highly scalable databases that support real time data collection, correlation and analysis, including (a) the Daytona database management system, (b) the GS data stream management system, (c) the Data Depot database generator and (d) the Bistro data feed manager. Together, these data systems have enabled the creation and maintenance of a database and data analysis infrastructure for troubleshooting complex issues in the network. This talk describes these data systems, presents their key research contributions, and identifies technical challenges that are currently being addressed.

Divesh Srivastava is the head of the Database Research Department at AT&T Labs-Research. He received his Ph.D. from the University of Wisconsin, Madison, and his B.Tech from the Indian Institute of Technology, Bombay. He is on the board of trustees of the VLDB Endowment, the associate Editor-in-Chief of the IEEE Transactions on Knowledge and Data Engineering, and an associate editor of the ACM Transactions on Database Systems. He has served as the program committee co-chair of many conferences, including VLDB 2007. His research interests and publications span a variety of topics in data management.

Keynote Talk 2: High-End Biological Imaging Generates Very Large 3D+ and Dynamic Datasets

Paul Matsudaira, Centre of Bioimaging Sciences, National University of Singapore

15th September (Wednesday), 9:00 AM – 10:30 AM

Location: Grand Ballroom

Biological imaging collects data over time and length scales that range from atoms and molecules to cells and tissues with a frequent goal of molecular detection in live tissues and organs. There are two generic approaches employed in microscopy to study the structure and function of molecules, cells, and organisms. Both collect large numbers of images and generate TB-size datasets but for different purposes. In high-content screening applications drugs are applied to live cells grown in wells and imaged in nautomated microscopes. A thorough study will examine the effects of a pharmacological agent at several dilutions by collecting images of cells at different time-points and generate TB datasets from millions of cells. After image processing and analysis, critical concentrations of drug activity are identified and mechanism of action is extracted from downstream information. The experiment is convoluted by marking various structures or molecules with a signature color, thus allowing for specific objects to be localized in space and tracked in time. The second approach has the goal of a 3D structure. Datasets consist of either serial slices or projections through the object, a cell, tissue, or organism and are acquired by light or electron microscopy. These methods are also linked in a time series to generate the dynamics of the structure from which mechanical and kinetic parameters are extracted. A significant problem is that these approaches can easily generate TBs of image data in minutes to hour periods. Thus, in biological imaging it is easy to generate data but the pressing problems are how to manage and organize image datasets, relate image datasets obtained by different microscopy methods, extract information from images, and to generate dynamic and 3D models of biological structures. These downstream steps contribute to the total dataset for the experiment.

Paul Matsudaira studied undergraduate biology and chemistry at the University of Washington (Seattle, WA) and received in 1981 a PhD in Biological Sciences from Dartmouth College (Hanover, NH) where he studied the structure of microvilli that line the intestine surface. In 1985, after postdoctoral studies in Germany (Max Planck for Biophysical Chemistry) and England (MRC Laboratory of Molecular Biology) he joined the faculty of the Biology Department at the Massachusetts Institute of Technology and the Whitehead Institute. After 23 years at MIT where he was Professor of Biology and Bioengineering, he moved in 2009 to the National University of Singapore as Professor and Head of the Department of Biological Sciences and founding Director of the NUS Centre for BioImaging Sciences. His lab studies the mechanics of cell movements and the structure of the cellular machinery that powers cell motility.

VLDB AWARDS

10 Year Best Paper Award Keynote 1: Dealing with Web Data: History and Look Ahead
Junghoo Cho, Hector Garcia-Molina
10 Year Best Paper: Evolution of the Web and Implications for an Incremental Crawler
16th September (Thursday), 9:00 AM – 9:45 AM
Location: Grand Ballroom

The high rate of change and the unprecedented scale of the Web pose enormous challenges to search engines who wish to provide the most up-to-date and highly relevant information to its users. The VLDB 2000 paper "The Evolution of the Web and Implications for an Incremental Crawler" tried to address part of this challenge by collecting and analyzing the Web history data and by describing the architecture and the associated algorithms for an incremental Web crawler that can provide more up-to-date data to users in a timely manner. Experiments and theoretical analysis showed - surprisingly at the time - that a policy that allocates more resources to more frequently changing items does not necessarily lead to better performance. In this paper, we discuss what has happened in the 10 years since and talk about the challenges that lie head.

Junghoo Cho *is an associate professor in the Department of Computer Science at University of California, Los Angeles. He received a Ph.D. degree in Computer Science from Stanford University and a B.S. degree in physics from Seoul National University. His main research interests are in the study of the evolution, management, retrieval and mining of information on the World-Wide Web. He publishes research papers in major international journals and conference proceedings. He serves on program committees of top international conferences, including SIGMOD, VLDB and WWW. He is a recipient of the NSF CAREER Award, IBM Faculty Award, Okawa Research Award and Northrop Grunmann Excellence in Teaching Award.*

Hector Garcia-Molina *is the Leonard Bosack and Sandra Lerner Professor in the Departments of Computer Science and Electrical Engineering at Stanford University, Stanford, California. He was the chairman of the Computer Science Department from January 2001 to December 2004. From 1997 to 2001 he was a member the President's Information Technology Advisory Committee (PITAC). From August 1994 to December 1997 he was the Director of the Computer Systems Laboratory at Stanford. From 1979 to 1991 he was on the faculty of the Computer Science Department at Princeton University, Princeton, New Jersey. His research interests include distributed computing systems, digital libraries and database systems. He received a BS in electrical engineering from the Instituto Tecnologico de Monterrey, Mexico, in 1974. From Stanford University, Stanford, California, he received in 1975 a MS in electrical engineering and a PhD in computer science in 1979. He holds an honorary PhD from ETH Zurich (2007).*

Garcia-Molina is a Fellow of the Association for Computing Machinery and of the American Academy of Arts and Sciences; is a member of the National Academy of Engineering; received the 1999 ACM SIGMOD Innovations Award; is a Venture Advisor for Onset Ventures, and is a member of the Board of Directors of Oracle.

10 Year Best Paper Award Keynote 2: Database Replication: A Tale of Research Across Communities

Bettina Kemme, Gustavo Alonso

10 Year Best Paper: Don't be lazy, be consistent: Postgres-R, a new way to implement database replication

16th September (Thursday), 9:45 AM – 10:30 AM

Location: Grand Ballroom

Replication has always been a key mechanism to achieve scalability and fault-tolerance in databases. In the last years its importance has been even further increased due to its role in providing elasticity at the database layer. In all these contexts, the biggest challenge lies in offering a replication solution that provides both performance and strong data consistency. Traditionally, performance could only be achieved through lazy replication at the expense of transactional guarantees. In contrast, eager, strong consistency approaches came with a performance penalty and poor scalability.

A decade ago, the way out of this situation involved combining results from distributed systems and databases. The use of group communication primitives with strong ordering and delivery guarantees combined with smart transaction handling (using snapshot isolation, transferring logs instead of re-executing updates, keeping the message overhead per transaction constant) were a radical departure from the state-of-the-art at the time. Now, these techniques are widely used in data centers and cloud computing. In this paper we review the context for the original work and discuss how these ideas have evolved and changed over the last 10 years.

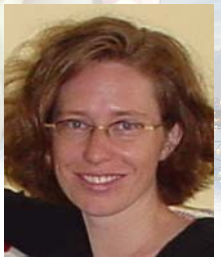

Bettina Kemme *is Associate Professor at the School of Computer Science of McGill University, Montreal, Canada. She received her undergraduate degree at the Friedrich-Alexander University in Erlangen, Germany, and her Ph.D. at the Swiss Federal Institute of Technology in Zurich (ETHZ). She also was a visiting student at UC Santa Barbara. Bettina's research focus lies in the design and devel opment of distributed information systems, in particular, in data consistency aspects and the interplay between communication and data management. Bettina has been PC member of many database and distributed systems conferences, such as VLDB, SIGMOD, ICDE, EDBT, Middleware, ICDCS, Eurosys, and P2P.*
She has been on the Editorial Board for the Encyclopedia of Database Systems, Springer, and track co-chair of ICDE 2009. She is area editor of Information Systems, Elsevier.

Gustavo Alonso *is a full professor at the Department of Computer Science (D-INFK) of the Swiss Federal Institute of Technology in Zurich (ETHZ) where he is a member of the Systems Group (www.systems.ethz.ch) and the Enterprise Computing Center (www.ecc.ethz.ch). He holds an engineering degree in telecommunications from the Madrid Technical University (ETSIT, Politecnica de Madrid, 1989) as well as an M.S. and a Ph.D. in Computer Science from UC Santa Barbara. Gustavo's main research interests include distributed systems, middleware, system aspects of software engineering, and data management. He has been general chair of MDM 2007, program chair or co-chair of Middleware 2005, VLDB 2006, BPM 2007, and ICDE 2008. He is in the Board of Trustees of the VLDB Endowment and the Chair of EuroSys, the European Chapter of ACM SIGOPS. He is also the CTO of Zimory GmbH.*

VLDB TUTORIALS

Tutorial 1: Distributed Caching Platforms
Anil Nori
14th September (Tuesday), 11:00 AM – 12:30 PM
Location: Grand Ballroom II

With the advances in processing, memory, and connectivity technologies, applications are becoming increasingly distributed, data-centric, and web based. These applications demand support for large number of users at high performance, requiring extreme scale, high availability, and low latency data access. Such applications are known as Extreme Transaction Processing (XTP) applications. Distributed Caching Platforms exploit the advances in the memory and networking technologies to fuse memory on multiple machines into a single unified global memory providing data access at low latencies. The Distributed Caching Platforms are evolving into in-memory data platforms for XTP applications. This tutorial provides a comprehensive overview of Distributed Caching Platform technologies, commercial products, and their usage.

Anil Nori *is a Distinguished Engineer in the Application Server Group (ASG) in Microsoft, focusing on application and data platform technologies. His current areas of focus include: Application Servers for Enterprise and Cloud deployments, Messaging Systems, Caching/In-memory databases, Rich client/mid-tier data and application architectures, Data/Application hosting, and Data/App manageability. He has over 30 years of experience in building complex database and application systems.*

Before coming to Microsoft, Anil was co-founder and CTO of Asera Inc., which built on-demand application platform, tools, and hosted solutions in Order Management and Supply Chain Management. Prior to Asera, Anil was at Oracle as a database server architect for Oracle7, Oracle8, and Oracle8i database releases. At Oracle, he drove the efforts in Oracle object-relational and extensible technology, Internet and multi-media DBMS development, and XML technologies. Before joining Oracle, Anil was a Database Architect for DEC database products, where he was involved in the development of centralized and distributed DBMS products. Anil also worked, as a Computer Scientist at Computer Corporation of America, on well-known distributed and semantic database research projects like SDD-1, ADAPLEX, and Multibase.

He is actively involved in the database and application server research community. He has published papers, gave talks, and presented tutorials at research and industry conferences. Anil holds over 25 patents in the field of database and application platform technologies.

Tutorial 2: Big Data and Cloud Computing: New Wine or Just New Bottles?
Divyakant Agrawal, Sudipto Das, Amr El Abbadi
14th September (Tuesday), 2:00 PM – 3:30 PM, 4:00 PM – 5:30 PM
Location: Grand Ballroom II

Cloud computing is an extremely successful paradigm of service oriented computing, and has revolutionized the way computing infrastructure is used. This success has seen a proliferation in the number of applications which leverage various cloud platforms, resulting in a tremendous increase in the scale of the data generated as well as consumed by such applications. Scalable database management systems – both for update intensive application workloads, as well as decision support systems for descriptive and deep analytics – are thus a critical part of the cloud infrastructure. The quest for conquering the challenges posed by the management of big data has led to a plethora of systems. This tutorial aims to clarify some of the critical concepts in the design space of big data and cloud computing such as: the appropriate systems for a specific set of application requirements, the research challenges in data management for the cloud, and what is novel in the cloud for database researchers? We also aim to address one basic question: Whether cloud computing poses new challenges in scalable data management or it is just a reincarnation of old problems? We provide a comprehensive background study of state-of-the-art systems for scalable data management and analysis. We also identify the critical aspects in the design of different systems and the applicability and scope of these systems. A thorough understanding of current solutions and a precise characterization of the design space are essential for clearing the "cloudy skies of data management" and ensuring the success of DBMSs in the cloud, thus emulating the success enjoyed by relational databases in traditional enterprise settings.

Dr. Divyakant Agrawal *is a Professor of Computer Science at the University of California at Santa Barbara. His research expertise is in the areas of database systems, distributed computing, data warehousing, and large-scale information systems. From January 2006 through December 2007, Dr. Agrawal served as VP of Data Solutions and Advertising Systems at the Internet Search Company ASK.com. While at ASK.com, Dr. Agrawal was the Chief Architect for building the next-generation Business Intelligence and Data Warehousing system. In addition, he developed revenue-sensitive products at ASK.com by applying data-mining and machine-learning technologies over ASK.com's historical data. Dr. Agrawal also served as a Visiting Senior Research Scientist at the NEC Laboratories of America in Cupertino, CA from 1997 to 2009. He is currently a Visiting Professor at the School of Computing at the National University of Singapore (Summer'2010). During his professional career, Dr. Agrawal has served extensively on the Program Committees of International Conferences, Symposia, and Workshops and served as an editor of the journal of Distributed and Parallel Databases (1993-2008), the VLDB journal (2003-2008) and currently serves on the editorial board of the Proceedings of the VLDB. He recently served as the Program Chair of the 2010 ACM International Conference on Management of Data and is currently serving as the General Chair of the 2010 ACM SIGSPATIAL Conference on Advances in Geographical Information Systems. Dr. Agrawal's research philosophy is to work on data management problem that have both practical as well as theoretical impact. He has published approximately 300 research manuscripts in prestigious forums (journals, conferences, symposia, and workshops) on wide range of topics related to data management and distributed systems.*

Sudipto Das r*eceived his B. Engg. degree in Computer Science and Engineering from Jadavpur University, India in 2006 and was awarded the TCS-JU best student award for 2006. He is currently a PhD candidate in the Department of Computer Science at University of California, Santa Barbara. His research interests lie in the area of scalable data management in cloud computing infrastructures. He is specifically interested in update intensive and transactional systems for cloud platforms that scale out while providing the developers consistency that they can reason about. He is the recipient of UCSB Computer Science Outstanding Teaching Assistant award for Fall 2008 and the UCSB graduate division Dissertation Fellowship for 2011. He is a student member of ACM and IEEE.*

Amr El Abbadi *received his PhD degree in Computer Science from Cornell University in 1987. He is currently a Professor and Chair of the Department of Computer Science at the University of California, Santa Barbara. He has held visiting professor positions at the University of Campinas in Brazil, IBM Almaden Research Center, the Swedish Institute of Computer Science in Stockholm, and at IRISA at the University of Rennes in France. He was Vice Chair of ICDCS 1999, Vice Chair for ICDE 2002, the Americas Program Chair for VLDB 2000, and is the Program Chair for SIGSPATIAL 2010 and Program co-chair of ACM SOCC 2011. He served as a board member of the VLDB Endowment from 2002 to 2008. In 2007, Prof. El Abbadi received the UCSB Senate Outstanding Mentorship Award for his excellence in mentoring graduate students. His research interests lie in the broad area of scalable database and distributed systems. He is a senior member of IEEE and ACM.*

Tutorial 3: Techniques for Similarity Searching in Multimedia Databases
Hanan Samet
15th September (Wednesday), 11:00 AM – 12:30 PM
Location: Grand Ballroom II

Similarity searching is a crucial part of retrieval in multimedia databases used for applications such as pattern recognition, image databases, and content-based retrieval. It involves finding objects in a data set S that are similar to a query object q based on some distance measure d which is usually a distance metric. The search process is usually achieved by means of nearest neighbor finding.
Existing methods for handling similarity search in this setting fall into one of two classes. The first is based on mapping to a vector space. The vector space is usually of high dimension which requires special handling due to the fact indexing methods do not discriminate well in such spaces. In particular, the query regions often overlap all of the blocks that result from the decomposition of the underlying space. This has led to some special solutions that make use of a sequential scan. An alternative is to use dimension reduction to find a mapping from a high-dimensional space into a low-dimensional space by finding the most discriminating dimensions and then index the data using one of a number of different data structures such as k-d trees, R-trees, quadtrees, etc. The second directly indexes the objects based on distances making use of data structures such as the vp-tree, M-tree, etc. At times, the distances are no t metrics which requires additional care.
This seminar is organized into four parts that include an overview as well as cover the basic concepts outlined above: indexing low and high dimensional spaces, distance-based indexing, and nearest neighbor searching.

Hanan Samet *(http://www.cs.umd.edu/~hjs/) is a Professor of Computer Science at the University of Maryland, College Park and is a member of the Institute for Computer Studies. He is also a member of the Computer Vision Laboratory at the Center for Automation Research where he leads a number of research projects on the use of hierarchical data structures for database applications involving spatial data. He has a Ph.D from Stanford University. He is the author of the recent book "Foundations of Multidimensional and Metric Data Structures" published by Morgan-Kaufmann, San Francisco, in 2006 (www.mkp.com/multidimensional), an award winner in the 2006 best book in Computer and Information Science competition of the Professional and Scholarly Publishers (PSP) Group of the American Publishers Association (AAP), and of the first two books on spatial data structures titled "Design and Analysis of Spatial Data Structures" and "Applications of Spatial Data Structures: Computer Graphics, Image Processing and GIS" published by Addison-Wesley, Reading, MA, 1990. He is the founding chair of ACM SIGSPATIAL, and a recipient of best paper awards in the 2008 SIGMOD Conference, the 2008 SIGSPATIAL Conference, and the 2007 Computers & Graphics Journal, the 2009 UCGIS Research Award and the 2010 CMPS Board of Visitors Award at the University of Maryland, a Fellow of the ACM, IEEE, and IAPR (International Association for Pattern Recognition), and an ACM Distinguished Speaker.*

Tutorial 4: Event Processing: Past, Present, Future
Opher Etzion
15th September (Wednesday), 2:00 PM – 3:30 PM, 4:00 PM – 5:00 PM
Location: Grand Ballroom II

Analysts have marked *Event Processing* as the most growing segment in enterprise computing during the last couple of years and this trend is expected to continue. Many the large and medium software companies (IBM, Oracle, Microsoft, Sybase, Progress Software and TIBCO) are now offering event processing platforms as well as 30 smaller companies. This area is now extensively coverage by analysts, establishment of a dedicated research community, books and academic courses.
The event processing commercial products are mostly descendents of research projects rooted in multiple disciplines, some of them are data management disciplines (active databases, data stream management, temporal databases) and some are rooted in other areas (discrete event simulation, distributed computing, formal verification).
The tutorial is intended for a technical audience that is interested in deep dive into understanding event processing. The audience will gain insights about event processing: What it really means? Where does it come from? How it relates to research concepts (e.g. stream computing) as well as enterprise computing terms (e.g. Business Rules Management Systems)? The audience will also gain insights into the current state of the art, the leading architectures, the basic building blocks of event processing, and the various programming styles exemplified by code examples. Last but not least, the audience will gain insights about the current trends, and the research challenges that exist, this part will be based on the discussions in the Event Processing Dagstuhl seminar that was held in May 2010.
This tutorial provides a deep dive on topics of: architecture, contexts and patterns of event processing. It explains the notions of event processing network and its different parts, and discusses implementation issues and optimization criteria and solutions and outline research challenges.

Dr. Opher Etzion *is IBM Senior Technical Staff Member; Event Processing Scientific Leader in IBM Haifa Research Lab, and chair of the EPTS (Event Processing Technical Society (en.wikipedia.org/wiki/Event_Processing_Technical_Society).*

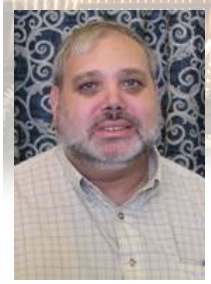

Previously he has been lead architect of event processing technology in IBM Websphere, and a Senior Manager in IBM Research division, managed a department that has performed one of the pioneering projects that shaped the area of "complex event processing". He is also blogging about event processing since August 1997 (http://epthinking.blogspot.com/).
He is a co-author (with Peter Niblett) of the first comprehensive technical book in the event processing area: Event Processing in Action (Manning, 2010) (http://www.manning.com/etzion/)
In parallel he serves as an adjunct professor at the Technion – Israel Institute of Technology, teaching an event processing course. He has authored or co-authored more than 70 papers in refereed journals and conferences, on topics related to: active databases, temporal databases, rule-base systems, complex event processing and autonomic computing, co-authored the book "Temporal Database – Research and Practice", Springer-Verlag, 1998, and supervised 25 PhD and MSc theses. Prior to joining IBM in 1997, he has been a faculty member and founding head of the Information Systems Engineering area at the Technion, and held professional and managerial positions in industry and in the Israel Air-Force.

Tutorial 5: Similarity Search and Mining in Uncertain Databases
Mathias Renz, Reynold Cheng, Hans-Peter Kriegel
16th September (Thursday), 11:00 AM – 12:30 PM
Location: Grand Ballroom II

Managing, searching and mining uncertain data has achieved much attention in the database community recently due to new sensor technologies and new ways of collecting data. There is a number of challenges in terms of collecting, modeling, representing, querying, indexing and mining uncertain data. In its scope, the diversity of approaches addressing these topics is very high because the underlying assumptions of uncertainty are different across different papers. This tutorial provides a comprehensive and comparative overview of general techniques for the key topics in the fields of querying, indexing and mining uncertain data. In particular, it identifies the most generic types of probabilistic similarity queries and discusses general algorithmic methods to answer such queries efficiently. In addition, the tutorial sketches probabilistic methods for important data mining applications in the context of uncertain data with special emphasis on probabilistic clustering and probabilistic pattern mining. The intended audience of this tutorial ranges from novice researchers to advanced experts as well as practitioners from any application domain dealing with uncertain data retrieval and mining.

Dr. Matthias Renz *is a Post-Doctoral Teaching and Research Assistant at the Institute for Informatics at the Ludwig-Maximilians University of Munich (LMU). He received his Diploma in Electrical Engineering in 1997 from the University of Applied Sciences (Munich) and Diploma in Computer Science in 2002 from LMU. In 2006, he received his PhD degree in Computer Science from the Ludwig-Maximilians University of Munich (LMU). Matthias Renz is a ACM SIGSPATIAL member and serves on the program committees and review panels for leading conferences, journals and workshops including VLDB, SSTD, CIKM, ACM GIS, WISE, ACM TODS, IEEE TKDE, VLDB Journal, Information Systems among others. 2006 he received the Best-Paper Award from the 11. International Conference on Database Systems for Advanced Applications (DASFAA'06). He was general chair of the 1st SIGSPATIAL ACM GIS 2009 International Workshop on Querying and Mining Uncertain Spatio-Temporal Data (QUeST'09). Matthias Renz was invited as keynote speaker for the 1st ACM SIGKDD Workshop on Knowledge Discovery from Uncertain Data (U'09). He presented tutorials at the SSTD'09 and DASFAA'09. His research interests include similarity search and mining in uncertain databases, traffic networks, spatial and temporal databases and sensor networks.*

Dr. Reynold Cheng *is the Assistant Professor of the Department of Computer Science in the University of Hong Kong (HKU). He received his BEng (Computer Engineering) in 1998, and MPhil (Computer Science and Information Systems) in 2000 from HKU. He then obtained his MSc and PhD degrees from the Department of Computer Science in Purdue University, in 2003 and 2005.*
Dr. Cheng was the Assistant Professor in the Department of Computing of the Hong Kong Polytechnic University from 2005 to 2008, where he received two Performance Awards. He is a member of IEEE, ACM, ACM SIGMOD, and UPE. He has served on the program committees and review panels for leading database conferences and journals like VLDB, ICDE, and TODS. He is also a guest editor for a special issue in TKDE. He is a keynote speaker in the First International Workshop on Quality of Context (QuaCon '09). His research interests include database management, as well as querying and mining of uncertain data.

Hans-Peter Kriegel *is a full professor for database systems and data mining in the Department "Institute for Informatics" at the Ludwig-Maximilians-Universität München, Germany and has served as the department chair or vice chair over the last years. His research interests are in spatial and multimedia database systems, particularly in query processing, performance issues, similarity search, high-dimensional indexing as well as in knowledge discovery and data mining. He has published over 300 refereed conference and journal papers and he received the "SIGMOD Best Paper Award" 1997 and the "DASFAA Best Paper Award" 2006 together with members of his research team. In 2009 he was appointed as an ACM fellow for contributions to knowledge discovery and data mining, similarity search, spatial data management, and access methods for high-dimensional data.*

Tutorial 6: Data Management and Mining in Internet Ad Systems
S Muthukrishnan
16th September (Thursday), 2:00 PM – 3:30 PM, 4:00 PM – 5:30 PM
Location: Grand Ballroom II

New systems produce data that often present new challenges for data management and mining problems. For example, inventory and sales data led to emphasis on data mining problems such as association rule mining; analysis of Internet Packet traffic logs led to data stream management systems; and, growing markup publication systems led to challenges addressed by semi-structured data management.

In this tutorial, we are inspired by systems that have emerged in the past decade that enable advertisements (ads) on the Internet. Such Internet ad systems handle billions of transactions every day involving millions of users, websites and advertisers, and are the basis for billions of dollars worth industry. They crucially rely on real-time collection, management and analysis of data for their effectiveness. Further, they represent unusual challenges for data analysis: nearly all parties in Internet ad systems from marketeers to publishers use active, selfish strategies that both help generate new data as well as distort data produced due to their selfish strategies. Mining such data while cognizant of the inherent game theory is a great research challenge. Finally, Internet ad systems use Information Retrieval, Auction and Game Theory, Machine Learning and Optimization Algorithms, and data analysis systems have to be compatible with these methods.

The tutorial will provide an overview of Internet ad systems and discuss in detail both data management as well as data mining tasks that arise: In the first part, we will describe different Internet ad systems and discuss issues in managing the data that arises in them as well as various tools. The second part will be on the data mining problems that arise, many unique to these systems.

The speaker ***S. (Muthu) Muthukrishnan*** *is a Professor in Rutgers University and a Research Scientist at Google. Muthu's research interest is in databases and algorithms, recently on data stream management and in algorithms for Internet ad systems.*


Notes:
VLDB2010

VLDB PANELS

Panel 1: Cloud Databases: What's New?
Moderator: Raghu Ramakrishnan, Yahoo! Research
Members: Daniel Abadi (Yale University, USA), Michael Carey (University of California, Irvine, USA), Surajit Chaudhuri (Microsoft Research, USA), Hector Garcia-Molina (Stanford University, USA), Jignesh M. Patel (University of Wisconsin, Madison, USA)
15th September (Wednesday), 2:00 PM – 3:30 PM
Location: Grand Ballroom I

The panelists will discuss what characterizes data management in the cloud, and how this differs from the broad range of applications that conventional database management systems have supported over the past few decades. They will examine whether we need to develop new technologies to address demonstrably new challenges, or whether we can largely re-position existing systems and approaches. The discussion will cover data analysis in the cloud using Map-Reduce based systems such as Hadoop, and cloud data serving (and so-called "No SQL" systems).

Raghu Ramakrishnan *is Chief Scientist for Audience & Cloud Computing, and a Fellow at Yahoo!, where he heads the Community Systems group. He has been Professor of Computer Sciences at the University of Wisconsin-Madison, and was founder and CTO of QUIQ, a company that pioneered question-answering communities, powering Ask Jeeves' AnswerPoint as well as customer-support for companies such as Compaq.*
His research is in the area of database systems, with a focus on data mining, online communities, and web-scale data management. He has developed scalable algorithms for clustering, decision-tree construction, and itemset counting, and was among the first to investigate mining of continuously evolving, stream data.
His work on query optimization and deductive databases has found its way into commercial database systems, and his work on extending SQL to deal with queries over sequences has influenced the design of window functions in SQL:1999. His paper on the Birch clustering algorithm received the SIGMOD 10-Year Test-of-Time award, and he has written the widely-used text "Database Management Systems" (WCB / McGraw-Hill, with J. Gehrke), now in its third edition.
He is Chair of ACM SIGMOD, on the Board of Directors of ACM SIGKDD and the Board of Trustees of the VLDB Endowment, and has served as editor-in-chief of the Journal of Data Mining and Knowledge Discovery, associate editor of ACM Transactions on Database Systems, and the Database area editor of the Journal of Logic Programming.
Raghu is a Fellow of the Association for Computing Machinery (ACM), and has received several awards, including a Distinguished Alumnus Award from IIT Madras, a Packard Foundation Fellowship, an NSF Presidential Young Investigator Award, and an ACM SIGMOD Contributions Award.

Panel 2: Time for Our Field to Grow Up
Moderator: Marianne Winslett, Univ. of Illinois, Urbana-Champaign, USA
Members: Anastassia Ailamaki (EPFL, Switzerland), Laura Haas (IBM Almaden Research Center), HV Jagadish (University of Michigan, USA), David Maier (Portland State University, USA), Tamer Ozsu (University of Waterloo, Canada)
16th September (Thursday), 4:00 PM – 5:30 PM
Location: Grand Ballroom I

Compared to centuries of physics and millennia of mathematics, the 50-year-history of computer science and information management research makes us the toddlers of the scientific community. Yet during our brief existence, we've revolutionized the world and, not content with that, gone on to build and study virtual worlds. We have justly taken pride in our accomplishments, and developed our own unique way of conducting research, unlike other scientific and engineering fields.

But cracks have appeared in this edifice we have built. The conference system that served us so well for our first 50 years is falling apart. Our ever-increasing population competes ever more energetically for a finite set of resources. Other scientific and engineering disciplines still think that our field equates to programming, and look down on us. While we may also look down on them, it is undeniably true that high-energy physicists get many more research dollars per capita than we do, and our computer science colleagues wonder whether all the data management problems haven't already been solved. Other departments have started to teach courses that overlap our turf. Are we our own worst enemies? Why doesn't everyone understand how important our research is? Do we have to abandon the conference system? Must we become more like the stodgy old fields of science and engineering? Or can we find our own way?

Marianne Winslett *is a professor in the Department of Computer Science at University of Illinois at Urbana-Champaign, and the director of the Advanced Digital Sciences Center in Singapore, which is owned by the University of Illinois. Her research interests lie in information security and in the management of scientific data. She is an ACM Fellow, and she received a Presidential Young Investigator Award from the National Science Foundation and two best paper awards for research on managing compliance data. She has served on the editorial boards of ACM Transactions on the Web, ACM Transactions on Database Systems, IEEE Transactions on Knowledge and Data Engineering, ACM Transactions on Information and Systems Security, and the Very Large Data Bases Journal. She is the former vice-chair of ACM SIGMOD.*

VLDB PAPER ABSTRACTS

Research Session 1: Database Security
Tuesday, 11:00 – 12:30

Building Disclosure Risk Aware Query Optimizers for Relational Databases
Mustafa Canim (University of Texas at Dallas, USA), Murat Kantarcioglu (University of Texas at Dallas, USA), Bijit Hore (University of California Irvine, USA), Sharad Mehrotra (University of California Irvine, USA)

Many DBMS products in the market provide built in encryption support to deal with the security concerns of the organizations. This solution is quite effective in preventing data leakage from compromised/stolen storage devices. However, recent studies show that a significant part of the leaked records have been done so by using specialized malwares that can access the main memory of systems. These malwares can easily capture the sensitive information that are decrypted in the memory including the cryptographic keys used to decrypt them. This can further compromise the security of data residing on disk that are encrypted with the same keys. In this paper we quantify the disclosure risk of encrypted data in a relational DBMS for main memory-based attacks and propose modifications to the standard query processing mechanism to minimize such risks. Specifically, we propose query optimization techniques and disclosure models to design a data-sensitivity aware query optimizer. We implemented a prototype DBMS by modifying both the storage engine and optimizer of MySQL-InnoDB server. The experimental results show that the disclosure risk of such attacks can be reduced dramatically while incurring a small performance overhead in most cases.

Secure Personal Data Servers: a Vision Paper
Tristan Allard (University of Versailles, France), Nicolas Anciaux (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, USA), Indrakshi Ray (Colorado State University, USA), Shaoyi Yin (Institut National de Recherche en Informatique et Automatique, France)

An increasing amount of personal data is automatically gathered and stored on servers by administrations, hospitals, insurance companies, etc. Citizen themselves often count on internet companies to store their data and make them reliable and highly available through the internet. However, these benefits must be weighed against privacy risks incurred by centralization. This paper suggests a radically different way of considering the management of personal data. It builds upon the emergence of new portable and secure devices combining the security of smart cards and the storage capacity of NAND Flash chips. By embedding a full-fledged Personal Data Server in such devices, user control of how her sensitive data is shared by others (by whom, for how long, according to which rule, for which purpose) can be fully reestablished and convincingly enforced. To give sense to this vision, Personal Data Servers must be able to interoperate with external servers and must provide traditional database services like durability, availability, query facilities, transactions. This paper proposes an initial design for the Personal Data Server approach, identifies the main technical challenges associated with it and sketches preliminary solutions. We expect that this paper will open exciting perspectives for future database research.

PolicyReplay: Misconfiguration-Response Queries for Data Breach Reporting
Daniel Fabbri (University of Michigan, USA), Kristen LeFevre (University of Michigan, USA), Qiang Zhu (University of Michigan, USA)

Recent legislation has increased the requirements of organizations to report data breaches, or unauthorized access to data. While access control policies are used to restrict access to a database, these policies are complex and difficult to configure. As a result, misconfigurations sometimes allow users access to unauthorized data.

In this paper, we consider the problem of reporting data breaches after such a misconfiguration is detected. To locate past SQL queries that may have revealed unauthorized information, we introduce the novel idea of a misconfiguration response (MR) query. The MR-query cleanly addresses the challenges of information propagation within the database by replaying the log of operations and returning all logged queries for which the result has changed due to the misconfiguration. A strawman implementation of the MR-query would go back in time and replay all the operations that occurred in the interim, with the correct policy. However, re-executing all operations is inefficient. Instead, we develop techniques to improve reporting efficiency by reducing the number of operations that must be re-executed and reducing the cost of replaying the operations. An extensive evaluation shows that our method can reduce the total runtime by up to an order of magnitude.

Research Session 2: Distributed and Parallel Databases
Tuesday, 11:00 – 12:30

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

Carlo Curino (Massachusetts Institute of Technology, USA), Yang Zhang (Massachusetts Institute of Technology, USA), Evan Jones (Massachusetts Institute of Technology, USA), Samuel Madden (Massachusetts Institute of Technology, USA)

We present Schism, a novel workload-aware approach for database partitioning and replication designed to improve scalability of shared-nothing distributed databases. Because distributed transactions are expensive in OLTP settings (a fact we demonstrate through a series of experiments), our partitioner attempts to minimize the number of distributed transactions, while producing balanced partitions. Schism consists of two phases: i) a workload-driven, graph-based replication/partitioning phase and ii) an explanation and validation phase. The first phase creates a graph with a node per tuple (or group of tuples) and edges between nodes accessed by the same transaction, and then uses a graph partitioner to split the graph into k balanced partitions that minimize the number of cross-partition transactions. The second phase exploits machine learning techniques to find a predicate-based explanation of the partitioning strategy (i.e., a set of range predicates that represent the same replication/ partitioning scheme produced by the partitioner).
The strengths of Schism are: i) independence from the schema layout, ii) effectiveness on n-to-n relations, typical in social network databases, iii) a unified and fine-grained approach to replication and partitioning. We implemented and tested a prototype of Schism on a wide spectrum of test cases, ranging from classical OLTP workloads (e.g., TPC-C and TPC-E), to more complex scenarios derived from social network websites (e.g., Epinions.com), whose schema contains multiple n-to-n relationships, which are known to be hard to partition. Schism consistently outperforms simple partitioning schemes, and in some cases proves superior to the best known manual partitioning, reducing the cost of distributed transactions up to 30%.

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)

Keyword search in relational databases has been extensively studied. Given a relational database, a keyword query finds a set of interconnected tuple structures connected by foreign key references. On RDBMS, a keyword query is processed in two steps, namely, candidate networks (CNs) generation and CNs evaluation, where a CN is an SQL. In common, a keyword query needs to be processed using over 10,000 SQLs. There are several approaches to process a keyword query on RDBMS, but there is a limit to achieve high performance on a uniprocessor architecture. In this paper, we study parallel computing keyword queries on a multicore architecture. We give three observations on keyword query computing, namely, a large number of SQLs that needs to be processed, high sharing possibility among SQLs, and large intermediate results with small number of final results. All make it challenging for parallel keyword queries computing. We investigate three approaches. We first study the query level parallelism, where each SQL is processed by one core. We distribute the SQLs into different cores based on three objectives, regarding minimizing workload skew, minimizing inter-core sharing and maximizing intra-core sharing respectively. Such an approach has the potential risk of load unbalancing through accumulating errors of cost estimation. We then study the operation level parallelism, where each operation of an SQL is processed by one core. All operations are processed in stages, where in each stage the costs of operations are re-estimated to reduce the accumulated error. Such operation level parallelism still has drawbacks of workload skew when large operations are involved and a large number of cores are used.

Finally, we propose a new algorithm that partitions relations adaptively in order to minimize the extra cost of partitioning and at the same time reduce workload skew. We conducted extensive performance studies using two large real datasets, DBLP and IMDB, and we report the efficiency of our approaches in this paper.

The Case for Determinism in Database Systems

Alexander Thomson (Yale University, USA), Daniel Abadi (Yale University, USA)

Replication is a widely used method for achieving high availability in database systems. Due to the nondeterminism inherent in traditional concurrency control schemes, however, special care must be taken to ensure that replicas don't diverge. Log shipping, eager commit protocols, and lazy synchronization protocols are well-understood methods for safely replicating databases, but each comes with its own cost in availability, performance, or consistency. In this paper, we propose a distributed database system which combines a simple deadlock avoidance technique with concurrency control schemes that guarantee equivalence to a predetermined serial ordering of transactions. This effectively removes all nondeterminism from typical OLTP workloads, allowing active replication with no synchronization overhead whatsoever. Further, our system eliminates the requirement for two-phase commit for any kind of distributed transaction, even across multiple nodes within the same replica. By eschewing deadlock detection and two-phase commit, our system under many workloads outperforms traditional systems that allow nondeterministic transaction reordering.

Research Session 3: Data Exchange
Tuesday, 11:00 – 12:30

MapMerge: Correlating Independent Schema Mappings

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

One of the main steps towards integration or exchange of data is to design the mappings that describe the (often complex) relationships between the source schemas or formats and the desired target schema. In this paper, we introduce a new operator, called MapMerge, that can be used to correlate multiple, independently designed schema mappings of smaller scope into larger schema mappings. This allows a more modular construction of complex mappings from various types of smaller mappings such as schema correspondences produced by a schema matcher or pre-existing mappings that were designed by either a human user or via mapping tools. In particular, the new operator also enables a new "divide-and-merge" paradigm for mapping creation, where the design is divided (on purpose) into smaller components that are easier to create and understand, and where MapMerge is used to automatically generate a meaningful overall mapping. We describe our MapMerge algorithm and demonstrate the feasibility of our implementation on several real and synthetic mapping scenarios. In our experiments, we make use of a novel similarity measure between two database instances with different schemas that quantifies the preservation of data associations. We show experimentally that MapMerge improves the quality of the schema mappings, by significantly increasing the similarity between the input source instance and the generated target instance.

Chase Termination: A Constraints Rewriting Approach

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

Several database areas such as data exchange and integration share the problem of fixing database instance violations with respect to a set of constraints. The chase algorithm solves such violations by inserting tuples and setting the value of nulls. Unfortunately, the chase algorithm may not terminate and the problem of deciding whether the chase process terminates is undecidable. Recently there has been an increasing interest in the identification of sufficient structural properties of constraints which guarantee that the chase algorithm terminates [8, 10, 14, 15].
In this paper we propose an original technique which allows to improve current conditions detecting chase termination. Our proposal consists in rewriting the original set of constraints Σ into an `equivalent' set Σ^{α} and verifying the structural properties for chase termination on Σ^{α}. The rewriting of constraints allows to recognize larger classes of constraints for which chase termination is guaranteed. In particular, we show that if Σ satisfies

chase termination conditions T , then the rewritten set $\sum^{\alpha}$ satisfies T as well, but the vice versa is not true, that is there are significant classes of constraints for which $\sum^{\alpha}$ satisfies T and $\sum$ does not.

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)

The recent literature has provided a solid theoretical foundation for the use of schema mappings in data-exchange applications. Following this formalization, new algorithms have been developed to generate optimal solutions for mapping scenarios in a highly scalable way, by relying on SQL. However, these algorithms suffer from a serious drawback: they are not able to handle key constraints and functional dependencies on the target, i.e., equality generating dependencies (egds). While egds play a crucial role in the generation of optimal solutions, handling them with first-order languages is a difficult problem. In fact, we start from a negative result: it is not always possible to compute solutions for scenarios with egds using an SQL script. Then, we identify many practical cases in which this is possible, and develop a best-effort algorithm to do this. Experimental results show that our algorithm produces solutions of better quality with respect to those produced by previous algorithms, and scales nicely to large databases.

Research Session 4: Database Services and Applications
Tuesday, 11:00 – 12:30

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 (ESRI, USA), Yehoshua Sagiv (Hebrew University of Jerusalem, Israel)

A route search is an enhancement of an ordinary geographic search. Instead of merely returning a set of entities, the result is a route that goes via entities that are relevant to the search. The input to the problem consists of several search queries, and each query defines a type of geographical entities. When visited, some of the entities succeed in satisfying the user while others fail to do so; however, only the probability of success is known prior to arrival. The main task is to find a route that visits at least one satisfying entity of each type. In an interactive search, the route is computed in steps. In each step, only the next entity of the route is given to the user, and after visiting that entity, the user provides a feedback specifying whether the entity satisfies her.
This paper investigates interactive route search in the presence of order constraints that specify that some types of entities should be visited before others. We present heuristic algorithms for interactive route search for two cases, depending on whether the constraints define a complete order or a partial one. The main challenge is to utilize the feedback in order to compute a route that is shorter and has a higher degree of success, compared to routes that are computed non-interactively. We also discuss how to compare the results of the algorithms and introduce suitable measures for doing so. Experiments on real-world data illustrate the efficiency and effectiveness of our algorithms.

Energy Management for MapReduce Clusters

Willis Lang (University of Wisconsin-Madison, USA), Jignesh Patel (University of Wisconsin-Madison, USA)

The area of cluster-level energy management has attracted significant research attention over the past few years. One class of techniques to reduce the energy consumption of clusters is to selectively power down nodes during periods of low utilization to increase energy efficiency. One can think of a number of ways of selectively powering down nodes, each with varying impact on the workload response time and overall energy consumption. Since the MapReduce framework is becoming "ubiquitous", the focus of this paper is on developing a framework for systematically considering various MapReduce node power down strategies, and their impact on the overall energy consumption and workload response time.
We closely examine two extreme techniques that can be accommodated in this framework. The first is based on a recently proposed technique called "Covering Set" (CS) that keeps only a small fraction of the nodes powered up during periods of low utilization. At the other extreme is a technique that we propose in this paper, called the All-In Strategy (AIS). AIS uses all the nodes in the cluster to run a workload and then powers down the

entire cluster. Using both actual evaluation and analytical modeling we bring out the differences between these two extreme techniques and show that AIS is often the right energy saving strategy.

Toward Scalable Keyword Search over Relational Data

Akanksha Baid (University of Wisconsin-Madison, USA), Ian Rae (University of Wisconsin-Madison, USA), Jiexing Li (University of Wisconsin-Madison, USA), AnHai Doan (University of Wisconsin-Madison, USA), Jeffrey Naughton (University of Wisconsin-Madison, USA)

Keyword search (KWS) over relational databases has recently received significant attention. Many solutions and many prototypes have been developed. This task requires addressing many issues, including robustness, accuracy, reliability, and privacy. An emerging issue, however, appears to be performance related: current KWS systems have unpredictable running times. In particular, for certain queries it takes too long to produce answers, and for others the system may even fail to return (e.g., after exhausting memory). In this paper we argue that as today's users have been "spoiled" by the performance of Internet search engines, KWS systems should return whatever answers they can produce quickly and then provide users with options for exploring any portion of the answer space not covered by these answers. Our basic idea is to produce answers that can be generated quickly as in today's KWS systems, then to show users query forms that characterize the unexplored portion of the answer space. Combining KWS systems with forms allows us to bypass the performance problems inherent to KWS without compromising query coverage. We provide a proof of concept for this proposed approach, and discuss the challenges encountered in building this hybrid system. Finally, we present experiments over real-world datasets to demonstrate the feasibility of the proposed solution.

Industry Session 1: Real-Time and Stream Processing

Tuesday, 11:00 – 12:30

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), Michalis Vlachos (IBM Research - Zurich, Switzerland)

The ever-increasing number of intrusions in public and commercial networks has created the need for high-speed archival solutions that continuously store streaming network data to enable forensic analysis and auditing. However, "turning back the clock" for post-attack analyses is not a trivial task. The first major challenge is that the solution has to sustain data archiving under extremely high-speed insertion rates. Moreover, the archives created need to be stored in a format that is compressed but still amenable to indexing. The above requirements make general-purpose databases unsuitable for this task, and, thus, dedicated solutions are required.

In this paper, we describe a prototype solution that satisfies all requirements for high-speed archival storage, indexing and data querying on network ow information. The superior performance of our approach is attributed to the on-the-fly compression and indexing scheme, which is based on compressed bitmap principles. Typical commercial solutions can currently process 20,000-60,000 ows per second. An evaluation of our prototype implementation on current commodity hardware using real-world traffic traces shows its ability to sustain insertion rates ranging from 500,000 to more than 1 million records per second. The system offers interactive query response times that enable administrators to perform complex analysis tasks on-the-fly. Our technique is directly amenable to parallel execution, allowing its application in domains that are challenged by large volumes of historical measurement data, such as network auditing, traffic behavior analysis and large-scale data visualization in service provider networks.

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, USA), Martin Hirzel (IBM Thomas J. Watson Research Center, USA), Henrique Andrade (IBM Thomas J. Watson Research Center, USA), Buğra Gedik (IBM Thomas J. Watson Research Center, USA), Kun-Lung Wu (IBM Thomas J. Watson Research Center, USA)

Applications from several domains are now being written to process live data originating from hardware and software-based streaming sources. Many of these applications have been written relying solely on database and data warehouse technologies, despite their lack of need for transactional support and ACID properties. In

several extreme high-load cases, this approach does not scale to the processing speeds that these applications demand. In this paper we demonstrate an application acceleration approach whereby a regular ODBC-based application is converted into a true streaming application with minimal disruption from a software engineering standpoint. We showcase our approach on three real-world applications. We experimentally demonstrate the substantial performance improvements that can be observed when contrasting the accelerated implementation with the original database-oriented implementation.

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 Chai (National University of Singapore, Republic of Singapore)

Currently, submission of new SNP entries into SNP repositories such as dbSNP by NCBI is done by manual curation. This gives rise to errors and ambiguities in SNP data entries. Due to the expo-nential increase in SNP discovery, there is a necessity to create algorithms to accurately and rapidly map SNPs as they are discov-ered in real time and depositing these entries automatically into a central SNP database. UASMAs are a set of algorithms to instan-taneously map SNPs efficiently and accurately by their unique chromosome position in real time. It is the result of integration of structures and algorithms in state of the art alignment methods MAQ, BWT-SW, Bowtie, SOAP2 and BWA.

Using BLAST employed by NCBI as benchmark where recall was at most 91%, recall performance of components Bowtie and BWA were much better at up to 99% for longer reads. Similarly, Bowtie and BWA performed better in terms of precision at greater than 91 % whereas BLAST was only 78 – 88%. BLAST performed poorly in terms of recall and precision for longer reads. Bowtie and BWA algorithms in UASMAs were superior in terms of perfor-mances in alignment of longer sequences and locating the precise chromosome position of any SNP with respect to the NCBI refer-ence assembly. Results obtained are fast, instantaneous and accurate.

Using UASMAs prove to be fast and optimal in mapping new variants onto the genome in view of depositing these entries accu-rately into a central database. Because it is done in real-time and with increased accuracy, recall and precision, the database created will be complete, up-to-date and devoid of ambiguities and redundancies.

Demo Sessions 1, 4: Core Database Technologies
Tuesday & Wednesday 11:00 – 12:30

The Picasso Database Query Optimizer Visualizer

Jayant Haritsa (Indian Institute of Science, India)

Abstract is not available in the paper.

CODS: Evolving Data Efficiently and Scalably in Column Oriented Databases

Ziyang Liu (Arizona State University, USA), Sivaramakrishnan Natarajan (Arizona State University, USA), Bin He (IBM Almaden Research Center, USA), Hui-I Hsiao (IBM Almaden Research Center, USA), Yi Chen (Arizona State University, USA)

Database evolution is the process of updating the schema of a database or data warehouse (schema evolution) and evolving the data to the updated schema (data evolution). Database evolution is often necessitated in relational databases due to the changes of data or workload, the suboptimal initial schema design, or the availability of new knowledge of the database. It involves two steps: updating the database schema, and evolving the data to the new schema. Despite the capability of commercial RDBMSs to well optimize query processing, evolving the data during a database evolution through SQL queries is shown to be prohibitively costly. We designed and developed CODS, a platform for efficient *data level* data evolution in column oriented databases, which evolves the data to the new schema without materializing query results or unnecessary compression/decompression as occurred in traditional *query level* approaches. CODS ameliorates the efficiency of data evolution by orders of magnitude compared with commercial or open source RDBMSs.

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)

In this demo, we present *fpga-ToPSS* (Toronto Publish/Subscribe System Family), an efficient event processing platform for highfrequency and low-latency algorithmic trading. Our event processing platform is built over reconfigurable hardware—FPGAs—to achieve line-rate processing. Furthermore, our event processing engine supports Boolean expression matching with an expressive predicate language that models complex financial strategies to autonomously buy and sell stocks based on real-time financial data.

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

Justin Levandoski (University of Minnesota, USA), Mohamed Mokbel (University of Minnesota, USA), Mohamed Khalefa (University of Minnesota, USA)

We demonstrate *CareDB*, a context and preference-aware database system. *CareDB* provides scalable *personalized* location-based services to users based on their preferences and current surrounding context. Unlike existing location-based database systems that answer queries based solely on proximity in distance, *CareDB* considers user preferences and various types of context in determining the answer to location-based queries. To this end, *CareDB* does not aim to define new location-based queries, instead, it aims to redefine the answer of existing location-based queries. To achieve its goals, *CareDB* has several distinguishing characteristics that revolve around a generic and extensible preference and context-aware query processing framework that addresses (a) scalable, efficient preference joins, (b) gracefully handling contextual attributes that are expensive to derive, and (c) support for uncertain attributes.

Cloudy: A Modular Cloud Storage System

Donald Kossmann (Eidgenössische Technische Hochschule Zürich, Switzerland), Tim Kraska (University of California Berkeley, USA), 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)

This demonstration presents Cloudy, a modular cloud storage system. Cloudy provides a highly exible architecture for distributed data storage and is designed to operate with multiple workloads. Based on a generic data model, Cloudy can be customized to meet application requirements. The goal of this demonstration is to show the ability of Cloudy to efficiently process different query languages, and to automatically adapt to varying load scenarios.

Geospatial Stream Query Processing using Microsoft SQL Server StreamInsight

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

Microsoft SQL Server spatial libraries contain several components that handle geometrical and geographical data types. With advances in geo-sensing technologies, there has been an increasing demand for geospatial streaming applications. Microsoft SQL Server StreamInsight (*StreamInsight*, for brevity) is a platform for developing and deploying streaming applications that run continuous queries over high-rate streaming events. With its extensibility infrastructure, StreamInsight enables developers to integrate their domain expertise within the query pipeline in the form of user defined modules.
This demo utilizes the extensibility infrastructure in Microsoft StreamInsight to leverage its continuous query processing capabilities in two directions. The first direction integrates SQL spatial libraries into the continuous query pipeline of StreamInsight. StreamInsight provides a well-defined temporal model over incoming events while SQL spatial libraries cover the spatial properties of events to deliver a solution for spatiotemporal stream query processing. The second direction extends the system with an *analytical refinement and prediction* layer. This layer analyzes historical data that has been accumulated and summarized over the years to refine, smooth and adjust the current query output as well as predict the output in the near future. The demo scenario is based on transportation data in Los Angeles County.

Using XMorph to Transform XML Data

Curtis Dyreson (Utah State University, USA), Sourav S Bhowmick (Nanyang Technological University, Republic of Singapore), Kirankanth Mallampalli (Utah State University, USA)

XMorph is a new, shape polymorphic, domain-specific XML query language. A query in a shape polymorphic language adapts to the shape of the input, freeing the user from having to know the input's shape and making the query applicable to a wide variety of differently shaped inputs. An XMorph query specifies the shape of the output. The XMorph query engine transforms the input to the desired shape by shredding an XML document to a graph of closest relationships, and performing a *closeness preserving* transformation. We plan to demonstrate XMorph using a Java applet, which can also be used by the audience during the demonstration, to evaluate various XMorph queries. The applet will show the output, the shapes generated by the query, and report on potential data loss in a transformation.

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

Di Wang (Worcester Polytechnic Institute, USA), Elke Rundensteiner (Worcester Polytechnic Institute, USA), Richard Ellison (University of Massachusetts Medical School, USA), Han Wang (Worcester Polytechnic Institute, USA)

Our analysis of many real-world event based applications has revealed that existing Complex Event Processing technology (CEP), while effective for efficient pattern matching on event stream, is limited in its capability of reacting in real-time to opportunities and risks detected or environmental changes. We are the first to tackle this problem by providing active rule support embedded directly within the CEP engine, henceforth called Active Complex Event Processing technology, or short, Active CEP. We design the Active CEP model and associated rule language that allows rules to be triggered by CEP system state changes and correctly executed during the continuous query process. Moreover we design an Active CEP infrastructure, that integrates the active rule component into the CEP kernel, allowing fine-grained and optimized rule processing. We demonstrate the power of Active CEP by applying it to the development of a collaborative project with UMass Medical School, which detects potential threads of infection and reminds healthcare workers to perform hygiene precautions in real-time.

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)

We demonstrate an efficient LINQ to SQL provider and its significant impact on the runtime performance of LINQ programs that process large data volumes. This alternative provider is based on Ferry, compilation technology that lets relational database systems participate in the evaluation of first-order functional programs over nested, ordered data structures. The Ferry-based provider seamlessly hooks into the .NET LINQ framework and generates SQL code that strictly adheres to the semantics of the LINQ data model. Ferry comes with strong code size guarantees and complete support for the LINQ Standard Query Operator family, enabling a truly interactive and compelling LINQ demonstration. A variety of inspection holes may be opened to learn about the internals of the Ferry-based LINQ to SQL provider.

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)

The workflow models have been essentially operation-centric for many years, ignoring almost completely the data aspects. Recently, a new paradigm of data-centric workflows, called business artifacts, has been introduced by Nigam and Caswell. We follow this approach and propose a model where artifacts are XML documents that evolve in time due to interactions with their environment, i.e. human users orWeb services. This paper proposes the AXART system as a distributed platform for collaborative work that harnesses the

power of our model. We will illustrate AXART with an example taken from the movie industry. Indeed, applying for a role in a film is a typical collaborative process that involves various participants, inside and outside the film company. The demonstration scenario considers both standard workflow process and dynamic workflow modifications, based on two extension mechanisms: workflow specialization and workflow exception. The workflows, modeled using artifacts, are supported by the AXART system by combining techniques specific to active documents, like view maintenance, with security techniques to manage access rights.

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, USA), Fan Ping (State University of New York University at Albany, USA), Jeong-Hyon Hwang (State University of New York University at Albany, USA)

We propose to demonstrate iFlow, our replication-based system that supports both fast and reliable processing of data streams over the Internet. iFlow uses a low degree of replication in conjunction with detouring techniques to overcome network outages. iFlow also deploys replicas in a manner that improves performance and availability at the same time, and can cope with varying system conditions by continually migrating replicas. Based on a live network monitoring application, our demonstration will substantiate the strengths of iFlow. During the demonstration, various visual tools will provide graphical evidence of improvements with regards to availability, performance, and resource usage. To show iFlow's adaptivity, these tools will also allow us to control the demonstration situation including injecting different types of failures.

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)

This is a demonstration of data coordination in a peer data management system through the employment of distributed triggers. The latter express in a declarative manner individual security and consistency requirements of peers, that cannot be ensured by default in the P2P environment. Peers achieve to handle in a transparent way data changes that come from local and remote actions and events. The distributed triggers are implemented as an extension of the active functionality of a centralized commercial DBMS. The language and execution semantics of distributed triggers are integrated in the kernel of the DBMS such that the latter handles transparently and simultaneously both centralized and distributed triggers. Moreover, the management of distributed triggers is associated with a set of peer acquaintance and termination protocols which are incorporated in the centralized DBMS.

Research Session 5: Data Models and Languages

Tuesday, 14:00 – 15:30

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

Barzan Mozafari (University of California Los Angeles, USA), Kai Zeng (University of California Los Angeles, USA), Carlo Zaniolo (University of California Los Angeles, USA)

There is growing interest in query language extensions for pattern matching over event streams and stored database sequences, due to the many important applications that such extensions make possible. The push for such extensions has led DBMS vendors and DSMS venture companies to propose Kleene-closure extensions of SQL standards, building on seminal research that demonstrated the effectiveness and amenability to efficient implementation of such constructs. These extensions, however powerful, suffer from limitations that severely impair their effectiveness in many real-world applications. To overcome these problems, we have designed the K*SQL language and system, based on our investigation of the *nested words*, which are recent models that generalize both words and trees.
K*SQL extends the existing relational sequence languages, and also enables applications from other domains such as genomics, software analysis, and XML processing. At the same time, K*SQL remains extremely efficient, using our powerful optimizations for pattern search over nested words. Furthermore, we show that

other sequence languages and XPath can be automatically translated into K*SQL, allowing for K*SQL to be also used as a high-performance query execution back-end for those languages. Therefore, K*SQL is a unifying SQL-based engine for sequence and XML queries, which provides novel optimization techniques for both.

Avalanche-Safe LINQ Compilation

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

We report on a query compilation technique that enables the construction of alternative e_cient query providers for Microsoft's Language Integrated Query (LINQ) framework. LINQ programs are mapped into an intermediate algebraic form, suitable for execution on any SQL:1999-capable relational database system.
This compilation technique leads to query providers that (1) faithfully preserve list order and nesting, both being core features of the LINQ data model, (2) support the complete family of LINQ's Standard Query Operators, (3) bring database support to LINQ to XML where the original provider performs in-memory query evaluation, and, most importantly, (4) emit SQL statement sequences whose size is only determined by the input query's result type (and thus independent of the database size).
A sample query scenario uses this LINQ provider to marry database-resident TPC-H and XMark data - resulting in a unique query experience that exhibits quite promising performance characteristics, especially for large data instances.

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), Wenyuan Yu (University of Edinburgh, United Kingdom)

A variety of integrity constraints have been studied for data cleaning. While these constraints can detect the presence of errors, they fall short of guiding us to correct the errors. Indeed, data repairing based on these constraints may not find *certain fixes* that are absolutely correct, and worse, may introduce new errors when repairing the data. We propose a method for finding certain fixes, based on master data, a notion of *certain regions*, and a class of *editing rules*. A certain region is a set of attributes that are assured correct by the users. Given a certain region and master data, editing rules tell us what attributes to fix and how to update them. We show how the method can be used in data monitoring and enrichment. We develop techniques for reasoning about editing rules, to decide whether they lead to a unique fix and whether they are able to fix all the attributes in a tuple, *relative to* master data and a certain region. We also provide an algorithm to identify minimal certain regions, such that a certain fix is warranted by editing rules and master data as long as one of the regions is correct. We experimentally verify the effectiveness and scalability of the algorithm.

Research Session 6: Semantics
Tuesday, 14:00 – 15:30

Explaining Missing Answers to SPJUA Queries

Melanie Herschel (Universität Tübingen, Germany), Mauricio Hernández (IBM Research, USA)

This paper addresses the problem of explaining missing answers in queries that include selection, projection, join, union, aggregation and grouping (SPJUA). Explaining missing answers of queries is useful in various scenarios, including query understanding and debugging. We present a general framework for the generation of these explanations based on source data. We describe the algorithms used to generate a correct, finite, and, when possible, minimal set of explanations. These algorithms are part of Artemis, a system that assists query developers in analyzing queries by, for instance, allowing them to ask why certain tuples are not in the query results. Experimental results demonstrate that Artemis generates explanations of missing tuples at a pace that allows developers to effectively use them for query analysis.

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, USA)

Violations of functional dependencies (FDs) are common in practice, often arising in the context of data integration or Web data extraction. Resolving these violations is known to be challenging for a variety of reasons, one of them being the exponential number of possible "repairs". Previous work has tackled this problem either by producing a single repair that is (nearly) optimal with respect to some metric, or by computing consistent answers to selected classes of queries without explicitly generating the repairs. In this paper, we propose a novel data cleaning approach that is not limited to finding a single repair or to a particular class of queries, namely, sampling from the space of possible repairs. We give several motivating scenarios where sampling from the space of FD repairs is desirable, propose a new class of useful repairs, and present an algorithm that randomly samples from this space. We also show how to restrict the space of generated repairs based on user-defined hard constraints that define an immutable trusted subset of the input relation, and we experimentally evaluate our algorithm against previous approaches. While this paper focuses on repairing FDs, we envision the proposed sampling approach to be applicable to other integrity constraints with large repair spaces.

Evaluating Entity Resolution Results

David Menestrina (Stanford University, USA), Steven Whang (Stanford University, USA), Hector Garcia-Molina (Stanford University, USA)

Entity Resolution (ER) is the process of identifying groups of records that refer to the same real-world entity. Various measures (e.g., pairwise F_1, cluster F_1) have been used for evaluating ER results. However, ER measures tend to be chosen in an ad-hoc fashion without careful thought as to what defines a good result for the specific application at hand. In this paper, our contributions are twofold. First, we conduct an analysis on existing ER measures, showing that they can often conict with each other by ranking the results of ER algorithms differently. Second, we explore a new distance measure for ER (called "generalized merge distance" or GMD) inspired by the edit distance of strings, using cluster splits and merges as its basic operations. A significant advantage of GMD is that the cost functions for splits and merges can be configured, enabling us to clearly understand the characteristics of a defined GMD measure. Surprisingly, a state-of-the-art clustering measure called Variation of Information is a special case of our configurable GMD measure, and the widely used pairwise F_1 measure can be directly computed using GMD. We present an efficient linear-time algorithm that correctly computes the GMD measure for a large class of cost functions that satisfy reasonable properties.

Research Session 7: Stream Databases

Tuesday, 14:00 – 15:30

High-Performance Dynamic Pattern Matching over Disordered Streams

Badrish Chandramouli (Microsoft Research, USA), Jonathan Goldstein (Microsoft Research, USA), David Maier (Portland State University, USA)

Current pattern-detection proposals for streaming data recognize the need to move beyond a simple regular-expression model over strictly ordered input. We continue in this direction, relaxing restrictions present in some models, removing the requirement for ordered input, and permitting stream revisions (modification of prior events). Further, recognizing that patterns of interest in modern applications may change frequently over the lifetime of a query, we support updating of a pattern specification without blocking input or restarting the operator. Our new pattern operator (called AFA) is a streaming adaptation of a non-deterministic finite automaton (NFA) where additional schema-based user-defined information, called a *register*, is accessible to NFA transitions during execution. AFAs support dynamic patterns, where the pattern itself can change over time. We propose clean order-agnostic patterndetection semantics for AFAs, with new algorithms that allow a very efficient implementation, while retaining significant expressiveness and supporting native handling of out-of-order input, stream revisions, dynamic patterns, and several optimizations. Experiments on Microsoft StreamInsight show that we achieve event rates of more than 200K events/sec (up to 5× better than simpler schemes). Our dynamic patterns give up to orders-of-magnitude better throughput than solutions such as operator restart, and our other optimizations are very effective, incurring low memory and latency.

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, USA), Renée Miller (University of Toronto, Canada), Nesime Tatbul (Eidgenössische Technische Hochschule Zürich, Switzerland)

There are many academic and commercial stream processing engines (SPEs) today, each of them with its own execution semantics. This variation may lead to seemingly inexplicable differences in query results. In this paper, we present SECRET, a model of the behavior of SPEs. SECRET is a descriptive model that allows users to analyze the behavior of systems and understand the results of window-based queries for a broad range of heterogeneous SPEs. The model is the result of extensive analysis and experimentation with several commercial and academic engines. In the paper, we describe the types of heterogeneity found in existing engines, and show with experiments on real systems that our model can explain the key differences in windowing behavior.

Recognizing Patterns in Streams with Imprecise Timestamps

Haopeng Zhang (University of Massachusetts Amherst, USA), Yanlei Diao (University of Massachusetts Amherst, USA), Neil Immerman (University of Massachusetts Amherst, USA)

Large-scale event systems are becoming increasingly popular in a variety of domains. Event pattern evaluation plays a key role in monitoring applications in these domains. Existing work on pattern evaluation, however, assumes that the occurrence time of each event is known precisely and the events from various sources can be merged into a single stream with a total or partial order. We observe that in real-world applications event occurrence times are often unknown or imprecise. Therefore, we propose a temporal model that assigns a time interval to each event to represent all of its possible occurrence times and revisit pattern evaluation under this model. In particular, we propose the formal semantics of such pattern evaluation, two evaluation frameworks, and algorithms and optimizations in these frameworks. Our evaluation results using both real traces and synthetic systems show that the event-based framework always outperforms the point-based framework and with optimizations, it achieves high efficiency for a wide range of workloads tested.

Research Session 8: Stream Databases
Tuesday, 14:00 – 15:30

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)

The RDF data model is gaining importance for applications in computational biology, knowledge sharing, and social communities. Recent work on RDF engines has focused on scalable performance for querying, and has largely disregarded updates. In addition to incremental bulk loading, applications also require online updates with exible control over multi-user isolation levels and data consistency. The challenge lies in meeting these requirements while retaining the capability for fast querying.
This paper presents a comprehensive solution that is based on an extended deferred-indexing method with integrated versioning. The version store enables time-travel queries that are efficiently processed without adversely affecting queries on the current data. For exible consistency, transactional concurrency control is provided with options for either snapshot isolation or full serializability. All methods are integrated in an extension of the RDF-3X system, and their very good performance for both queries and updates is demonstrated by measurements of multi-user workloads with real-life data as well as stress-test synthetic loads.

Graph Pattern Matching: From Intractable to Polynomial Time

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (National University of Defense 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)

Graph pattern matching is typically defined in terms of subgraph isomorphism, which makes it an NP-complete problem. Moreover, it requires bijective functions, which are often too restrictive to characterize patterns in emerging applications. We propose a class of graph patterns, in which an edge denotes the connectivity in a data graph within a predefined number of hops. In addition, we define matching based on a notion of bounded simulation, an extension of graph simulation. We show that with this revision, graph pattern matching can be performed in cubic-time, by providing such an algorithm. We also develop algorithms for incrementally finding matches when data graphs are updated, with performance guarantees for DAG patterns. We experimentally verify that these algorithms scale well, and that the revised notion of graph pattern matching allows us to identify communities commonly found in real-world networks.

GRAIL: Scalable Reachability Index for Large Graphs

Hilmi Yildirim (Rensselaer Polytechnic Institute, USA), Vineet Chaoji (Yahoo! Research Labs, USA), Mohammed Zaki (Rensselaer Polytechnic Institute, USA)

Given a large directed graph, rapidly answering reachability queries between source and target nodes is an important problem. Existing methods for reachability trade-off indexing time and space versus query time performance. However, the biggest limitation of existing methods is that they simply do not scale to very large real-world graphs. We present a very simple, but scalable reachability index, called GRAIL, that is based on the idea of randomized interval labeling, and that can effectively handle very large graphs. Based on an extensive set of experiments, we show that while more sophisticated methods work better on small graphs, GRAIL is the only index that can scale to millions of nodes and edges. GRAIL has linear indexing time and space, and the query time ranges from constant time to being linear in the graph order and size.

Research Session 9: Middleware Platforms for Data Management
Tuesday, 14:00 – 15:30

HaLoop: Efficient Iterative Data Processing on Large Clusters

Yingyi Bu (University of Washington, USA), Bill Howe (University of Washington, USA), Magdalena Balazinska (University of Washington, USA), Michael Ernst (University of Washington, USA)

The growing demand for large-scale data mining and data analysis applications has led both industry and academia to design new types of highly scalable data-intensive computing platforms. MapReduce and Dryad are two popular platforms in which the dataflow takes the form of a directed acyclic graph of operators. These platforms lack built-in support for iterative programs, which arise naturally in many applications including data mining, web ranking, graph analysis, model fitting, and so on. This paper presents HaLoop, a modified version of the Hadoop MapReduce framework that is designed to serve these applications. HaLoop not only extends MapReduce with programming support for iterative applications, it also dramatically improves their efficiency by making the task scheduler loop-aware and by adding various caching mechanisms. We evaluated HaLoop on real queries and real datasets. Compared with Hadoop, on average, HaLoop reduces query runtimes by 1.85, and shuffles only 4% of the data between mappers and reducers.

The Impact of Virtual Views on Containment

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

Virtual views are a mechanism that facilitates re-use and makes queries easier to express. However the use of iterative view definitions makes very simple query evaluation and analysis problems more complex. In this paper we study classical containment and equivalence problems for queries built up through simple unions of conjunctive queries and view definitions. More precisely, we determine the complexity of containment and equivalence for non-recursive Datalog. We show that the problem is much harder than its classical counterpart - complete for co-NEXPTIME. We then show that this remains true even with restrictions on the schema and queries in place. Finally, we isolate subcases that are more tractable, ranging from NP to PSPACE.

Updatable and Evolvable Transforms for Virtual Databases

James Terwilliger (Microsoft Corporation, USA), Lois Delcambre (Portland State University, USA), David Maier (Portland State University, USA), Jeremy Steinhauer (Portland State University, USA), Scott Britell (Portland State University, USA)

Applications typically have some local understanding of a database schema, a *virtual database* that may differ significantly from the actual schema of the data where it is stored. Application engineers often support a virtual database using custom-built middleware because the available solutions, including updatable views, are unable to express necessary capabilities. We propose an alternative means of mapping a virtual database to a physical database that guarantees they remain synchronized under data or schema updates against the virtual schema. One constructs a mapping by composing *channel transformations* (CTs) that encapsulate atomic transformations — including complex transformations such as pivoting — with known updatability properties. Applications, query interfaces, and any other services can behave as if the virtual database is the implemented schema. We describe how CTs translate queries, DML, and DDL, and the properties that are necessary for such translation to be correct. We describe two example CTs in detail, and evaluate an implementation of channels for completeness and performance.

Demo Sessions 2, 5: Web Data Management and Information Retrieval Systems

Tuesday & Wednesday 14:00 – 15:30

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, USA), Lizhu Zhou (Tsinghua University, People's Republic of China)

Form-style interfaces have been widely used to allow users to access information. In this demonstration paper, we develop a new search paradigm in form-style query interfaces, called *Seaform* (which stands for Search-As-You-Type in Forms), which computes answers on-the-fly as a user types in a query letter by letter and gives the user instant feedback. Seaform provides better user experiences compared with traditional form-based query systems by reducing the efforts for a user to compose a high-quality query to find relevant answers. Seaform can also enhance faceted search and allow users to on-the-fly explore the underlying data. This search paradigm requires high performance to achieve an interactive speed. We develop efficient techniques and use them to implement two systems on real datasets. We demonstrate the features of these systems.

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

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

Spatial and temporal data have become ubiquitous in many application domains such as the Geosciences or life sciences. Sophisticated database management systems are employed to manage such structured data. However, an important source of spatio-temporal information that has not been fully utilized are unstructured text documents. In documents, combinations of temporal and spatial expressions form events, which can be mapped to a database structure and organized into trajectories that can be explored. In this context, the coupling of information retrieval techniques with spatio-temporal database concepts leads to new ways for managing and exploring document collections.
In this demonstration, we present TimeTrails, a system for the extraction, querying, storage, and exploration of spatiotemporal information embedded in text documents. The user can query a document collection, and TimeTrails visualizes the spatio-temporal information extracted from relevant documents as document trajectories, resulting in a map-based view of documents. This view helps the user to explore the temporal and spatial content of documents in a meaningful way and to further restrict search results using spatial and temporal predicates.

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)

Abstract is not available in the paper.

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)

Database systems often use XML schema to describe the format of valid XML documents. Usually, this format is determined when the system is designed. Sometimes, in an already functioning system, a need arises to change the XML schemas. In such a situation, the system has to transform the old XML documents so that they conform to the new format and that as little information as possible is lost in the process. This process is called *schema evolution*.
We have implemented an XML schema transformation toolkit within IBM Master Data Management Server (MDM). MDM uses XML documents to describe products that an enterprise may be offering to its clients. In this work we focus on evolving schemas rather than on integrating separate or heterogeneous data sources. Our solution includes an extendible schema matching algorithm that was designed with evolving XML schemas in mind and takes advantage of hierarchical structure of XML. It also includes a data transformation and migration method appropriate for environments where migration is performed in an abstraction layer above the DBMS. Finally, we describe a novel way of extending an XSLT editor with an XSLT visualization feature to allow the user's input and evaluation of the transformation.

XSACT: A Comparison Tool for Structured Search Results

Ziyang Liu (Arizona State University, USA), Sivaramakrishnan Natarajan (Arizona State University, USA), Peng Sun (Arizona State University, USA), Stephen Booher (Arizona State University, USA), Tim Moohan (Arizona State University, USA), Robert Winkler (Arizona State University, USA), Yi Chen (Arizona State University, USA)

Studies show that about 50% of web search is for information exploration purpose, where a user would like to investigate, compare, evaluate, and synthesize multiple relevant results. Due to the absence of general tools that can effectively analyze and differentiate multiple results, a user has to manually read and comprehend potentially large results in an exploratory search. Such a process is time consuming, labor intensive and error prone. With meta information embedded, keyword search on structured data provides the potential for automating or semi-automating the comparison of multiple results.
In this demo we present a system XSACT for differentiating search results on structured data. XSACT takes as input a set of structured results, and outputs a Differentiation Feature Set (DFS) for each result to highlight their differences within a size bound. The problem of generating DFSs with maximal differences is proved to be NP-hard. XSACT adopts efficient algorithms for DFS generation, and features a user-friendly interface that effectively interacts with the users to help them compare search results.

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)

We present in this paper ObjectRunner, a system for extracting, integrating and querying structured data from the Web. Our system harvests real-world items from template-based HTML pages (the so-called structured Web). It illustrates a two-phase querying of the Web, in which an intentional description of the targeted data is first provided, in a flexible and widely applicable manner. ObjectRunner follows then a lightweight, best-effort approach, leveraging both the input description and the source structure. This process is domain-independent, in the sense that it applies to any relation, either flat or nested, describing real-world items. We advocate via our prototype that fully automatic extraction and integration of structured data can be done fast and effectively, when the redundancy of the Web meets knowledge over the to-be-extracted data. We present the technical details and the overall platform through several application scenarios on real-life Web sources.

ROXXI: Reviving witness dOcuments to eXplore 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)

In recent years, there has been considerable research on information extraction and constructing RDF knowledge bases. In general, the goal is to extract all relevant information from a corpus of documents, store it into an ontology, and answer future queries based only on the created knowledge base. Thus, the original documents become dispensable. On the one hand, an ontology is a convenient and non-redundant structured source of information, based on which specific queries can be answered efficiently. On the other hand, many users doubt the correctness of facts and ontology subgraphs presented to them as query results without proof. Instead, users often wish to verify the obtained facts or subgraphs by reading about them in context, i.e., in a document relating the facts and providing background information. In this demo, we present ROXXI, a system operating on top of an existing knowledge base and reviving the abandoned witness documents. In doing so, it goes the opposite way of information extraction approaches – starting with ontological facts and tracing their way back to the documents they were extracted from. ROXXI offers interfaces for expert users (SPARQL) as well as for non-experts (ontology browser) and provides a ranked list of documents each associated with a content snippet highlighting the queried facts in context. At the demonstration site, we will show the advantages of this novel approach towards document retrieval and illustrate the benefits of reviving the documents that information extraction approaches neglect.

EXTRUCT: Using Deep Structural Information in XML Keyword Search

Arash Termehchy (University of Illinois at Urbana-Champaign, USA), Marianne Winslett (University of Illinois at Urbana-Champaign, USA)

Users who are unfamiliar with database query languages can search XML data sets using keyword queries. Previous work has shown that current XML keyword search methods, although intuitive, do not effectively use the data's structural information and provide poor precision, recall, and ranking for most queries. Based on an extension of the concept of information theory, we have developed principled frameworks called normalized total correlation (NTC) and normalized term presence correlation (NTPC) to measure the relevance of candidate answers to keyword queries. We demonstrate EXTRUCT, an XML keyword search interface that uses NTC and NTPC. An extensive empirical evaluation over two real-world XML DBs has shown that EXTRUCT has better precision and recall and provides better ranking than all previous approaches. We demonstrate EXTRUCT, along with seven other keyword search systems for four real-world XML data sets, using prepared queries as well as queries from the audience. The demonstration shows that using deep structural information increases the effectiveness of XML keyword search systems considerably.

SQL QueRIE Recommendations

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

This demonstration presents QueRIE, a recommender system that supports interactive database exploration. This system aims at assisting non-expert users of scientific databases by tracking their querying behavior and generating personalized query recommendations. The system is supported by two recommendation engines and the underlying recommendation algorithms. The first identifies potentially "interesting" parts of the database related to the corresponding data analysis task by locating those database parts that were accessed by similar users in the past. The second identifies structurally similar queries to the ones posted by the current user. Both approaches result in a recommendation set of SQL queries that is provided to the user to modify, or directly post to the database. The demonstrated system will enable users to query and get real-time recommendations from the SkyServer database, using user traces collected from the SkyServer query log.

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)

As the amount of user generated content grows, personal information management has become a challenging problem. Several information management approaches, such as desktop search, document organization and (collaborative) document tagging have been proposed to address this, however they are either inappropriate or inefficient. Automated collaborative document tagging approaches mitigate the problems of manual tagging, but they are usually based on centralized settings which are plagued by problems such as scalability, privacy, etc. To resolve these issues, we present P2PDocTagger, an automated and distributed document tagging system based on classification in P2P networks. P2PDocTagger minimizes the efforts of individual peers and reduces computation and communication cost while providing high tagging accuracy, and eases of document organization/retrieval. In addition, we provide a realistic and flexible simulation toolkit – P2PDMT, to facilitate the development and testing of P2P data mining algorithms.

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)

Web archives are useful resources to find out about the temporal evolution of persons, organizations, products, or other topics. However, even when advanced text search functionality is available, gaining insights into the temporal evolution of a topic can be a tedious task and often requires sifting through many documents. The demonstrated system named INZEIT (pronounced "insight") assists users by determining insightful time points for a given query. These are the time points at which the top-k time-travel query result changes substantially and for which the user should therefore inspect query results. INZEIT determines the m most insightful time points efficiently using an extended segment tree for in-memory bookkeeping.

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

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

Tags associated with social images are valuable information source for superior image search and retrieval experiences. Due to the nature of tagging, many tags associated with images are not visually descriptive. Consequently, presence of these noisy tags may reduce the effectiveness of tags' role in image retrieval. To address this problem, we demonstrate iAVATAR (interActive VisuAl-representative TAgs Relationship) system that uses the notion of *Normalized Image Tag Clarity* (NITC) to find *visual-representative tags*. A visual-representative tag effectively describes the visual content of the images. Further, we visually demonstrate relationships between *popular* tags and visual-representative tags as well as *co-occurrence* likelihood of a pair of tags associated with a *search tag* or image using *tag relationship graph* (TRG). We demonstrate various innovative features of iAVATAR with a real-world dataset and show that it enriches users' understanding of various important tag features during image search.

Research Session 10: Novel/Advanced Applications
Tuesday, 16:00 – 17:30

Navigating in Complex Mashed-Up Applications

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

Mashups integrate a set of Web-services and data sources, often referred to as mashlets. We study in this paper a common scenario where these mashlets are components of larger Web-Applications. In this case,

integration of mashlets yields a set of inter-connected applications, referred to as Mashed-up Applications (abbr. MashAPP). While interactions between the mashlets enrich the individual applications, they also render navigation within them more intricate for the user, as actions in one application may affect others. To assist users in their navigation through MashAPPs we provide a solution based on a simple, generic model for MashAPPs and navigation flows within them. Queries over the model allow users to describe navigation flows of interest, and an effective query evaluation algorithm provides users with recommendations on how to navigate within the MashAPP. The model and algorithms serve as a basis for the COMPASS system, built on top of the Mashup Server.

Dremel: Interactive Analysis of Web-Scale Datasets

Sergey Melnik (Google, USA), Andrey Gubarev (Google, USA), Jing Jing Long (Google, USA), Geoffrey Romer (Google, USA), Shiva Shivakumar (Google, USA), Matt Tolton (Google, USA), Theo Vassilakis (Google, USA)

Dremel is a scalable, interactive ad-hoc query system for analysis of read-only nested data. By combining multi-level execution trees and columnar data layout, it is capable of running aggregation queries over trillion-row tables in seconds. The system scales to thousands of CPUs and petabytes of data, and has thousands of users at Google. In this paper, we describe the architecture and implementation of Dremel, and explain how it complements MapReduce-based computing. We present a novel columnar storage representation for nested records and discuss experiments on few-thousand node instances of the system.

On Graph Query Optimization in Large Networks

Peixiang Zhao (University of Illinois at Urbana-Champaign, USA), Jiawei Han (University of Illinois at Urbana-Champaign, USA)

The dramatic proliferation of sophisticated networks has resulted in a growing need for supporting effective querying and mining methods over such large-scale graph-structured data. At the core of many advanced network operations lies a common and critical graph query primitive: *how to search graph structures efficiently within a large network?* Unfortunately, the graph query is hard due to the NP-complete nature of subgraph isomorphism. It becomes even challenging when the network examined is large and diverse. In this paper, we present a high performance graph indexing mechanism, SPath, to address the graph query problem on large networks. SPath leverages decomposed shortest paths around vertex neighborhood as basic indexing units, which prove to be both effective in graph search space pruning and highly scalable in index construction and deployment. Via SPath, a graph query is processed and optimized beyond the traditional vertex-at-a-time fashion to a more efficient *path-at-a-time* way: the query is first decomposed to a set of shortest paths, among which a subset of candidates with good selectivity is picked by a query plan optimizer; Candidate paths are further joined together to help recover the query graph to finalize the graph query processing. We evaluate SPath with the state-of-the-art GraphQL on both real and synthetic data sets. Our experimental studies demonstrate the effectiveness and scalability of SPath, which proves to be a more practical and efficient indexing method in addressing graph queries on large networks.

Research Session 11: Ranking Queries

Tuesday, 16:00 – 17:30

Proximity Rank Join

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

We introduce the proximity rank join problem, where we are given a set of relations whose tuples are equipped with a score and a real-valued feature vector. Given a target feature vector, the goal is to return the K combinations of tuples with high scores that are as close as possible to the target and to each other, according to some notion of distance. The setting closely resembles that of traditional rank join, but the geometry of the vector space plays a distinctive role in the computation of the overall score of a combination. Also, the input relations typically return their results either by distance from the target or by score. Because of these aspects, it turns out that traditional rank join algorithms, such as the well-known HRJN, have shortcomings in solving the proximity rank join problem, as they may read more input than needed. To overcome this weakness, we define a tight bound (used as a stopping criterion) that guarantees instance optimality, i.e., an I/O cost is achieved that is always within a constant factor of optimal. The tight bound can also be used to drive an adaptive pulling

strategy, deciding at each step which relation to access next. For practically relevant classes of problems, we show how to compute the tight bound efficiently. An extensive experimental study validates our results and demonstrates significant gains over existing solutions.

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

Akrivi Vlachou (Norwegian University of Science and Technology, Norway), Christos Doulkeridis (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)

Top-k queries are widely applied for retrieving a ranked set of the k most interesting objects based on the individual user preferences. As an example, in online marketplaces, customers (users) typically seek a ranked set of products (objects) that satisfy their needs. Reversing top-k queries leads to a query type that instead returns the set of customers that find a product appealing (it belongs to the top-k result set of their preferences). In this paper, we address the challenging problem of processing queries that identify the top-m most *influential products* to customers, where influence is defined as the cardinality of the reverse top-k result set. This definition of influence is useful for market analysis, since it is directly related to the number of customers that value a particular product and, consequently, to its visibility and impact in the market. Existing techniques require processing a reverse top-k query for each object in the database, which is prohibitively expensive even for databases of moderate size. In contrast, we propose two algorithms, SB and BB, for identifying the most influential objects: SB restricts the candidate set of objects that need to be examined, while BB is a branch-and-bound algorithm that retrieves the result incrementally. Furthermore, we propose meaningful variations of the query for most influential objects that are supported by our algorithms. Our experiments demonstrate the efficiency of our algorithms both for synthetic and real-life datasets.

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)

The location-aware keyword query returns ranked objects that are near a query location and that have textual descriptions that match query keywords. This query occurs inherently in many types of mobile and traditional web services and applications, e.g., Yellow Pages and Maps services. Previous work considers the potential results of such a query as being independent when ranking them. However, a relevant result object with nearby objects that are also relevant to the query is likely to be preferable over a relevant object without relevant nearby objects.

The paper proposes the concept of prestige-based relevance to capture both the textual relevance of an object to a query and the effects of nearby objects. Based on this, a new type of query, the Location-aware top-*k* Prestige-based Text retrieval (L*k*PT) query, is proposed that retrieves the top-*k* spatial web objects ranked according to both prestige-based relevance and location proximity. We propose two algorithms that compute L*k*PT queries. Empirical studies with real-world spatial data demonstrate that L*k*PT queries are more effective in retrieving web objects than a previous approach that does not consider the effects of nearby objects; and they show that the proposed algorithms are scalable and outperform a baseline approach significantly.

Research Session 12: Spatial and Temporal Databases
Tuesday, 16:00 – 17:30

Parsimonious Linear Fingerprinting for Time Series

Lei Li (Carnegie Mellon University, USA), B. Aditya Prakash (Carnegie Mellon University, USA), Christos Faloutsos (Carnegie Mellon University, USA)

We study the problem of mining and summarizing multiple time series effectively and efficiently. We propose PLiF, a novel method to discover essential characteristics ("fingerprints"), by exploiting the joint dynamics in numerical sequences. Our fingerprinting method has the following benefits: (a) it leads to *interpretable* features; (b) it is *versatile*: PLiF enables numerous mining tasks, including clustering, compression, visualization, forecasting, and segmentation, matching top competitors in each task; and (c) it is fast and *scalable*, with linear complexity on the length of the sequences.

We did experiments on both synthetic and real datasets, including human motion capture data (17MB of human motions), sensor data (166 sensors), and network router traffic data (18 million raw updates over 2 years). Despite its generality, PLiF outperforms the top clustering methods on clustering; the top compression methods on compression (3 times better reconstruction error, for the same compression ratio); it gives meaningful visualization and at the same time, enjoys a *linear* scale-up.

The HV-tree: a Memory Hierarchy Aware Version Index

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

The huge amount of temporal data generated from many important applications call for a highly efficient and scalable version index. The TSB-tree has the potential of large scalability due to its unique feature of progressive migration of data to larger mediums. However, its traditional design optimized for two levels of the memory hierarchy (the main memory and the hard disk) undermines its potential for high efficiency in face of today's advances in hardware, especially CPU/cache speed and memory size. We propose a novel version index structure called the HV-tree. Different from all previous version index structures, the HV-tree has nodes of different sizes, each optimized for a level of the memory hierarchy. As data migrates to different levels of the memory hierarchy, the HV-tree will adjust the node size automatically to exploit the best performance of all levels of the memory hierarchy. Moreover, the HV-tree has a unique chain mechanism to maximally keep recent data in higher levels of the memory hierarchy. As a result, HV-tree is several times faster than the TSB-tree for point queries (query with single key and single time value), and up to 1000 times faster than the TSB-tree for key-range and time-range queries.

Transforming Range Queries To Equivalent Box Queries To Optimize Page Access

Sakti Pramanik (Michigan State University, USA), Alok Watve (Michigan State University, USA), Chad Meiners (Michigan State University, USA), Alex Liu (Michigan State University, USA)

Range queries based on L_1 distance are a common type of queries in multimedia databases containing feature vectors. We propose a novel approach that transforms the feature space into a new feature space such that range queries in the original space are mapped into equivalent box queries in the transformed space. Since box queries are axes aligned, there are several implementational advantages that can be exploited to speed up the retrieval of query results. For two dimensional data the transformation is precise. For greater than two dimensions we propose a space transformation scheme based on disjoint planer rotation, and along with pruning query box the results are precise. Experimental results with large synthetic databases and some real databases show the effectiveness of the proposed transformation scheme. These experimental results have been corroborated with appropriate mathematical models.

Research Session 13: Record Linkage

Tuesday, 16:00 – 17:30

Record Linkage with Uniqueness Constraints and Erroneous Values

Songtao Guo (AT&T Interactive, USA), Xin Dong (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA), Remi Zajac (AT&T Interactive, USA)

Many data-management applications require integrating data from a variety of sources, where different sources may refer to the same real-world entity in different ways and some may even provide erroneous data. An important task in this process is to recognize and merge the various references that refer to the same entity. In practice, some attributes satisfy a *uniqueness* constraint—each real-world entity (or most entities) has a unique value for the attribute (*e.g.*, business contact phone, address, and email). Traditional techniques tackle this case by first linking records that are likely to refer to the same real-world entity, and then fusing the linked records and resolving conflicts if any. Such methods can fall short for three reasons: first, erroneous values from sources may prevent correct linking; second, the real world may contain exceptions to the uniqueness constraints and always enforcing uniqueness can miss correct values; third, locally resolving conflicts for linked records may overlook important global evidence.

This paper proposes a novel technique to solve this problem. The key component of our solution is to reduce the problem into a *k*-partite graph clustering problem and consider in clustering both similarity of attribute values and the sources that associate a pair of values in the same record. Thus, we perform global linkage and

fusion simultaneously, and can identify incorrect values and differentiate them from alternative representations of the correct value from the beginning. In addition, we extend our algorithm to be tolerant to a few violations of the uniqueness constraints. Experimental results show accuracy and scalability of our technique.

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)

Entity linkage is central to almost every data integration and data cleaning scenario. Traditional techniques use some computed similarity among data structure to perform merges and then answer queries on the merged data. We describe a novel framework for entity linkage with uncertainty. Instead of using the linkage information to merge structures a-priori, possible linkages are stored alongside the data with their belief value. A new probabilistic query answering technique is used to take the probabilistic linkage into consideration. The framework introduces a series of novelties: (i) it performs merges at run time based not only on existing linkages but also on the given query; (ii) it allows results that may contain structures not explicitly represented in the data, but generated as a result of a reasoning on the linkages; and (iii) enables an evaluation of the query conditions that spans across linked structures, offering a functionality not currently supported by any traditional probabilistic databases. We formally define the semantics, describe an efficient implementation and report on the findings of our experimental evaluation.

Behavior Based Record Linkage

Mohamed Yakout (Purdue University, USA), Ahmed Elmagarmid (Purdue University, USA), Hazem Elmeleegy (Purdue University, USA), Mourad Ouzzani (Purdue University, USA), Alan Qi (Purdue University, USA)

In this paper, we present a new record linkage approach that uses *entity behavior* to decide if potentially different entities are in fact the same. An entity's behavior is extracted from a transaction log that records the *actions* of this entity with respect to a given data source. The core of our approach is a technique that merges the behavior of two possible matched entities and computes the *gain* in recognizing behavior patterns as their matching score. The idea is that if we obtain a well recognized behavior after merge, then most likely, the original two behaviors belong to the same entity as the behavior becomes more complete after the merge. We present the necessary algorithms to model entities' behavior and compute a matching score for them. To improve the computational efficiency of our approach, we precede the actual matching phase with a fast candidate generation that uses a "quick and dirty" matching method. Extensive experiments on real data show that our approach can significantly enhance record linkage quality while being practical for large transaction logs.

Research Session 14: Experimental Analysis and Performance
Tuesday, 16:00 – 17:30

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)

Graphs are of growing importance in modeling complex structures such as chemical compounds, proteins, images, and program dependence. Given a query graph Q, the *subgraph isomorphism* problem is to find a set of graphs containing Q from a graph database, which is NP-complete. Recently, there have been a lot of research efforts to solve the subgraph isomorphism problem for a large graph database by utilizing graph indexes. By using a graph index as a filter, we prune graphs that are not real answers at an inexpensive cost. Then, we need to use expensive subgraph isomorphism tests to verify filtered candidates only. This way, the number of disk I/Os and subgraph isomorphism tests can be significantly minimized. The current practice for experiments in graph indexing techniques is that the author of a newly proposed technique does not implement existing indexes on his own code base, but instead uses the original authors' binary executables and reports only the wall clock time. However, we observe this practice may result in several problems. In order to address these problems, we have made significant efforts in implementing all representative indexing methods on a common framework called iGraph. Unlike existing implementations which either use (full or partial) in-memory

representations or rely on OS file system cache without guaranteeing real disk I/Os, we have implemented these indexes on top of a storage engine that guarantees real disk I/Os. Through extensive experiments using many synthetic and real datasets, we also provide new empirical findings in the performance of the full disk-based implementations of these methods.

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)

One of the main reasons why cloud computing has gained so much popularity is due to its ease of use and its ability to scale computing resources on demand. As a result, users can now rent computing nodes on large commercial clusters through several vendors, such as Amazon and rackspace. However, despite the attention paid by Cloud providers, performance unpredictability is a major issue in Cloud computing for (1) database researchers performing wall clock experiments, and (2) database applications providing servicelevel agreements. In this paper, we carry out a study of the performance variance of the most widely used Cloud infrastructure (Amazon EC2) from different perspectives. We use established microbenchmarks to measure performance variance in CPU, I/O, and network. And, we use a multi-node MapReduce application to quantify the impact on real dataintensive applications. We collected data for an entire month and compare it with the results obtained on a local cluster. Our results show that EC2 performance varies a lot and often falls into two bands having a large performance gap in-between — which is somewhat surprising. We observe in our experiments that these two bands correspond to the different virtual system types provided by Amazon. Moreover, we analyze results considering different availability zones, points in time, and locations. This analysis indicates that, among others, the choice of availability zone also influences the performance variability. A major conclusion of our work is that the variance on EC2 is currently so high that wall clock experiments may only be performed with considerable care. To this end, we provide some hints to users.

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)

MapReduce has been widely used for large-scale data analysis in the Cloud. The system is well recognized for its elastic scalability and fine-grained fault tolerance although its performance has been noted to be suboptimal in the database context. According to a recent study [19], Hadoop, an open source implementation of MapReduce, is slower than two state-of-the-art parallel database systems in performing a variety of analytical tasks by a factor of 3.1 to 6.5. MapReduce can achieve better performance with the allocation of more compute nodes from the cloud to speed up computation; however, this approach of "renting more nodes" is not cost effective in a pay-as-you-go environment. Users desire an economical elastically scalable data processing system, and therefore, are interested in whether MapReduce can offer both elastic scalability and efficiency.
In this paper, we conduct a performance study of MapReduce (Hadoop) on a 100-node cluster of Amazon EC2 with various levels of parallelism. We identify five design factors that affect the performance of Hadoop, and investigate alternative but known methods for each factor. We show that by carefully tuning these factors, the overall performance of Hadoop can be improved by a factor of 2.5 to 3.5 for the same benchmark used in [19], and is thus more comparable to that of parallel database systems. Our results show that it is therefore possible to build a cloud data processing system that is both elastically scalable and efficient.

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)

Despite the huge amount of recent research efforts on entity resolution (matching) there has not yet been a comparative evaluation on the relative effectiveness and efficiency of alternate approaches. We therefore present such an evaluation of existing implementations on challenging real-world match tasks. We consider approaches both with and without using machine learning to find suitable parameterization and combination of similarity functions. In addition to approaches from the research community we also consider a state-of-the-art commercial entity resolution implementation. Our results indicate significant quality and efficiency differences

between different approaches. We also find that some challenging resolution tasks such as matching product entities from online shops are not sufficiently solved with conventional approaches based on the similarity of attribute values.

Demo Sessions 3, 6: Data Extraction, Integration and Mining
Tuesday & Wednesday 16:00 – 17:30

Deep Web Integration with VisQI
Thomas Kabisch (Humboldt Universität zu Berlin, Germany), Eduard Dragut (University of Illinois at Chicago, USA), Clement Yu (University of Illinois at Chicago, USA), Ulf Leser, Humboldt Universität zu Berlin, Germany)

In this paper, we present VisQI (VISual Query interface Integration system), a Deep Web integration system. VisQI is capable of (1) transforming Web query interfaces into hierarchically structured representations, (2) of classifying them into application domains and (3) of matching the elements of different interfaces. Thus VisQI contains solutions for the major challenges in building Deep Web integration systems. The system comes along with a full-edged evaluation system that automatically compares generated data structures against a gold standard. VisQI has a framework-like architecture such that other developers can reuse its components easily.

SOLOMON: Seeking the Truth Via Copying Detection
Xin Dong (AT&T Labs - Research, USA), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

We live in the Information Era, with access to a huge amount of information from a variety of data sources. However, data sources are of different qualities, often providing conflicting, out-of-date and incomplete data. Data sources can also easily copy, reformat and modify data from other sources, propagating erroneous data. These issues make the identification of high quality information and sources non-trivial.
We demonstrate the SOLOMON system, whose core is a module that detects copying between sources. We demonstrate that we can effectively detect copying relationship between data sources, leverage the results in truth discovery, and provide a user-friendly interface to facilitate users in identifying sources that best suit their information needs.

Just-in-time Data Integration in Action
Martin Hentschel (Eidgenössische Technische Hochschule Zürich, Switzerland), Laura Haas (IBM Almaden Research Center, USA), Renée Miller (University of Toronto, Canada)

Today's data integration systems must be flexible enough to support the typical iterative and incremental process of integration, and may need to scale to hundreds of data sources. In this work we present a novel data integration system that offers great flexibility and scalability. Our approach to data integration is unique in that it executes mapping rules at query runtime using annotations. On top, we have built the *People People People* application. It allows users to search for people, display information about people, and browse through a network of related people, where the data is integrated from local and remote data sources. The demo presents all features of our underlying data integration engine through a set of motivating scenarios.

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 Heimel (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)

Abstract is not available in the paper.

Using Sentinel Technology in the TARGIT BI Suite

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

This paper demonstrates so-called *sentinels* in the TARGIT BI Suite. Sentinels are a novel type of rules that can warn a user if one or more *measure changes* in a multi-dimensional data cube are expected to cause a change to another measure critical to the user. We present the concept of sentinels, and we explain how sentinels represent *stronger* and more *specific* rules than sequential patterns and correlation techniques. In addition, we present the algorithm, implementation, and data warehouse setup that are prerequisites for our demo. In the demo we present a dialogue where users, without any prior technical knowledge, are able to select a critical measure, a number of cubes, and a time dimension, and subsequently mine and schedule sentinels for *early warnings*.

CoDA: Interactive Cluster Based Concept Discovery

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

Large data resources are ubiquitous in science and business. For these domains, an intuitive view on the data is essential to fully exploit the hidden knowledge. Often, these data can be semantically structured by concepts. Since the determination of concepts requires a thorough analysis of the data, data mining methods have to be applied. In the field of subspace clustering, some techniques have recently shown to be effective for this task. Although these methods generate concept-based patterns, the user has to provide domain knowledge to gain reasonable concepts out of the data.
Our demonstration CoDA (Concept Determination and Analysis) is a tool that supports the user in the final step of concept definition. More concretely, the user is guided through an iterative, interactive process in which concepts are suggested, analyzed, and potentially refined. The core aspect of CoDA is an intuitive, concept-driven presentation of subspace clusters such that concepts can be visually captured.

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 Velegrakis (University of Trento, Italy)

We propose the demonstration of Keymantic, a system for keyword-based searching in relational databases that does not require a-priori knowledge of instances held in a database. It finds numerous applications in situations where traditional keyword-based searching techniques are inapplicable due to the unavailability of the database contents for the construction of the required indexes.

Data Auditor: Exploring Data Quality and Semantics using Pattern Tableaux

Lukasz Golab (AT&T Labs - Research, USA), Howard Karloff (AT&T Labs - Research, USA), Flip Korn (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

We present Data Auditor, a tool for exploring data quality and data semantics. Given a rule or an integrity constraint and a target relation, Data Auditor computes pattern tableaux, which concisely summarize subsets of the relation that (mostly) satisfy or (mostly) fail the constraint. This paper describes 1) the architecture and user interface of Data Auditor, 2) the supported constraints for testing data consistency and completeness, 3) the heuristics used by Data Auditor to "tune" a given constraint or its associated parameters for better fit with the data, and 4) several demonstration scenarios using real data sets.

Research Session 15: Cloud Computing
Wednesday, 11:00 – 12:30

MRShare: Sharing Across Multiple Queries in MapReduce

Tomasz Nykiel (University of Toronto, Canada), Michalis Potamias (Boston University, USA), Chaitanya Mishra (Facebook, USA), George Kollios (Boston University, USA), Nick Koudas (University of Toronto, Canada)

Large-scale data analysis lies in the core of modern enterprises and scientific research. With the emergence of cloud computing, the use of an analytical query processing infrastructure (e.g., Amazon EC2) can be directly mapped to monetary value. MapReduce has been a popular framework in the context of cloud computing, designed to serve long running queries (jobs) which can be processed in batch mode. Taking into account that different jobs often perform similar work, there are many opportunities for sharing. In principle, sharing similar work reduces the overall amount of work, which can lead to reducing monetary charges incurred while utilizing the processing infrastructure. In this paper we propose a sharing framework tailored to MapReduce.
Our framework, MRShare, transforms a batch of queries into a new batch that will be executed more efficiently, by merging jobs into groups and evaluating each group as a single query. Based on our cost model for MapReduce, we define an optimization problem and we provide a solution that derives the optimal grouping of queries. Experiments in our prototype, built on top of Hadoop, demonstrate the overall effectiveness of our approach and substantial savings.

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)

Cloud storage is an emerging infrastructure that offers Platforms as a Service (PaaS). On such platforms, storage and compute power are adjusted dynamically, and therefore it is important to build a highly scalable and reliable storage that can elastically scale on demand with minimal startup cost.
In this paper, we propose ecStore – an elastic cloud storage system that supports automated data partitioning and replication, load balancing, efficient range query, and transactional access. In ecStore, data objects are distributed and replicated in a cluster of commodity computer nodes located in the cloud. Users can access data via transactions which bundle read and write operations on multiple data items stored on possibly different cluster nodes.
The architecture of ecStore follows a stratum design that leverages an underlying distributed index with a replication layer in the middle and a transaction management layer on top. ecStore provides adaptive read consistency on replicated data. We also enhance the system with an effective load balancing scheme using a self-tuning replication technique that is specially designed for large-scale data. Furthermore, a multi-version optimistic concurrency control scheme matches well with the characteristics of data in cloud storages. To validate the performance of the system, we have conducted extensive experiments on various platforms including a commercial cloud (Amazon's EC2), an in-house cluster, and PlanetLab.

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 (Max-Planck Institute for Informatics, Germany), Jörg Schad (Saarland University, Germany)

MapReduce is a computing paradigm that has gained a lot of attention in recent years from industry and research. Unlike parallel DBMSs, MapReduce allows non-expert users to run complex analytical tasks over very large data sets on very large clusters and clouds. However, this comes at a price: MapReduce processes tasks in a scan-oriented fashion. Hence, the performance of Hadoop—an open-source implementation of MapReduce—often does not match the one of a well-configured parallel DBMS. In this paper we propose a new type of system named Hadoop++: it boosts task performance without changing the Hadoop framework at all (Hadoop does not even 'notice it'). To reach this goal, rather than changing a working system (Hadoop), we inject our technology at the right places through UDFs only and affect Hadoop from inside. This has three important consequences: First, Hadoop++ significantly outperforms Hadoop. Second, any future changes of Hadoop may directly be used with Hadoop++ without rewriting any glue code. Third, Hadoop++ does not need to change

the Hadoop interface. Our experiments show the superiority of Hadoop++ over both Hadoop and HadoopDB for tasks related to indexing and join processing.

Research Session 16: Query Processing I
Wednesday, 11:00 – 12:30

Slicing Long-Running Queries
Nicolas Bruno (Microsoft Research, USA), Vivek Narasayya (Microsoft Research, USA), Ravi Ramamurthy (Microsoft Research, USA)

The ability to decompose a complex, long-running query into simpler queries that produce the same result is useful for many scenarios, such as admission control, resource management, fault tolerance, and load balancing. In this paper we propose query slicing as a novel mechanism to do such decomposition. We study different ways to extend a traditional query optimizer to enable query slicing and experimentally evaluate the benefits of each approach.

Sharing-Aware Horizontal Partitioning for Exploiting Correlations during Query Processing
Kostas Tzoumas (Aalborg University, Denmark), Amol Deshpande (University of Maryland College Park, USA), Christian Jensen (Aarhus University, Denmark)

Optimization of join queries based on average selectivities is suboptimal in highly correlated databases. In such databases, relations are naturally divided into partitions, each partition having substantially different statistical characteristics. It is very compelling to discover such data partitions during query optimization and create multiple plans for a given query, one plan being optimal for a particular combination of data partitions. This scenario calls for the sharing of state among plans, so that common intermediate results are not recomputed. We study this problem in a setting with a routing-based query execution engine based on eddies [1]. Eddies naturally encapsulate horizontal partitioning and maximal state sharing across multiple plans. We define the notion of a *conditional join plan*, a novel representation of the search space that enables us to address the problem in a principled way. We present a low-overhead greedy algorithm that uses statistical summaries based on *graphical models*. Experimental results suggest an order of magnitude faster execution time over traditional optimization for high correlations, while maintaining the same performance for low correlations.

Advanced Processing for Ontological Queries
Andrea Cali (University of Oxford, United Kingdom), Georg Gottlob (University of Oxford, United Kingdom), Andreas Pieris (Brunei University, United Kingdom)

Ontology-based data access is a powerful form of extending database technology, where a classical extensional database (EDB) is enhanced by an ontology that generates new intensional knowledge which may contribute to answer a query. The ontological integrity constraints for generating this intensional knowledge can be specified in description logics such as DL-Lite. It was recently shown that these formalisms allow for very efficient query-answering. They are, however, too weak to express simple and useful integrity constraints that involve joins. In this paper we introduce a more expressive formalism that takes joins into account, while still enjoying the same low query-answering complexity. In our framework, ontological constraints are expressed by sets of rules that are so-called tuple-generating dependencies (TGDs). We propose the language of sticky sets of TGDs, which are sets of TGDs with a restriction on multiple occurrences of variables (including joins) in the rule bodies. We establish complexity results for answering conjunctive queries under sticky sets of TGDs, showing, in particular, that ontological conjunctive queries can be compiled into first-order and thus SQL queries over the given EDB instance. We also show how sticky sets of TGDs can be combined with functional dependencies. In summary, we obtain a highly expressive and effective ontological modeling language that unifies and generalizes both classical database constraints and important features of the most widespread tractable description logics.

Research Session 17: Data Extraction
Wednesday, 11:00 – 12:30

Towards The Web of Concepts: Extracting Concepts from Large Datasets
Aditya Parameswaran (Stanford University, USA), Hector Garcia-Molina (Stanford University, USA), Anand Rajaraman (Kosmix Corporation, USA)

Concepts are sequences of words that represent real or imaginary entities or ideas that users are interested in. As a first step towards building a web of concepts that will form the backbone of the next generation of search technology, we develop a novel technique to extract concepts from large datasets. We approach the problem of concept extraction from corpora as a market-basket problem, adapting statistical measures of support and confidence. We evaluate our concept extraction algorithm on datasets containing data from a large number of users (e.g., the AOL query log data set), and we show that a high-precision concept set can be extracted.

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)

We propose a novel extraction approach that exploits *content redundancy* on the web to extract structured data from *template-based* web sites. We start by populating a seed database with records extracted from a few initial sites. We then identify values within the pages of each new site that match attribute values contained in the seed set of records. To match attribute values with diverse representations across sites, we define a new similarity metric that leverages the templatized structure of attribute content. Specifically, our metric discovers the matching pattern between attribute values from two sites, and uses this to ignore extraneous portions of attribute values when computing similarity scores. Further, to filter out noisy attribute value matches, we exploit the fact that attribute values occur at fixed positions within template-based sites. We develop an efficient Apriori-style algorithm to systematically enumerate attribute position configurations with sufficient matching values across pages. Finally, we conduct an extensive experimental study with real-life web data to demonstrate the effectiveness of our extraction approach.

Automatic Rule Refinement for Information Extraction
Bin Liu (University of Michigan, USA), Laura Chiticariu (IBM Almaden Research Center, USA), Vivian Chu (IBM Almaden Research Center, USA), H. Jagadish (University of Michigan, USA), Frederick Reiss (IBM Almaden Research Center, USA)

Rule-based information extraction from text is increasingly being used to populate databases and to support structured queries on unstructured text. Specification of suitable information extraction rules requires considerable skill and standard practice is to refine rules iteratively, with substantial effort. In this paper, we show that techniques developed in the context of data provenance, to determine the lineage of a tuple in a database, can be leveraged to assist in rule refinement. Specifically, given a set of extraction rules and correct and incorrect extracted data, we have developed a technique to suggest a ranked list of rule modifications that an expert rule specifier can consider. We implemented our technique in the *SystemT* information extraction system developed at IBM Research – Almaden and experimentally demonstrate its effectiveness.

Research Session 18: Privacy
Wednesday, 11:00 – 12:30

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)

Users of text search engines are increasingly wary that their activities may disclose confidential information about their business or personal profiles. It would be desirable for a search engine to perform document retrieval for users while protecting their intent. In this paper, we identify the privacy risks arising from semantically related search terms within a query, and from recurring high-specificity query terms in a search session. To counter the risks, we propose a solution for a similarity text retrieval system to offer anonymity and plausible deniability for the query terms, and hence the user intent, without degrading the system's precision-recall performance. The solution comprises a mechanism that embellishes each user query with decoy terms that exhibit similar specificity spread as the genuine terms, but point to plausible alternative topics. We also provide an accompanying retrieval scheme that enables the search engine to compute the encrypted document relevance scores from only the genuine search terms, yet remain oblivious to their distinction from the decoys. Empirical evaluation results are presented to substantiate the effectiveness of our solution.

Small Domain Randomization: Same Privacy, More Utility

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

Random perturbation is a promising technique for privacy preserving data mining. It retains an original sensitive value with a certain probability and replaces it with a random value from the domain with the remaining probability. If the replacing value is chosen from a large domain, the retention probability must be small to protect privacy. For this reason, previous randomization-based approaches have poor utility. In this paper, we propose an alternative way to randomize sensitive values, called *small domain randomization*. First, we partition the given table into sub-tables that have smaller domains of sensitive values. Then, we randomize the sensitive values within each sub-table independently. Since each sub-table has a smaller domain, a larger retention probability is permitted. We propose this approach as an alternative to classical partition-based approaches to privacy preserving data publishing. There are two key issues: ensure the published sub-tables do not disclose more private information than what is permitted on the original table, and partition the table so that utility is maximized. We present an effective solution.

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, USA), Dimitris Papadias (The Hong Kong University of Science and Technology, People's Republic of China)

The tremendous growth of the Internet has significantly reduced the cost of obtaining and sharing information about individuals, raising many concerns about user privacy. Spatial queries pose an additional threat to privacy because the location of a query may be sufficient to reveal sensitive information about the querier. In this paper we focus on k nearest neighbor (kNN) queries and define the notion of strong location privacy, which renders a query indistinguishable from any location in the data space. We argue that previous work fails to support this property for arbitrary kNN search. Towards this end, we introduce methods that offer strong location privacy, by integrating private information retrieval (PIR) functionality. Specifically, we employ secure hardware-aided PIR, which has been proven very efficient and is currently considered as a practical mechanism for PIR. Initially, we devise a benchmark solution building upon an existing PIR-based technique. Subsequently, we identify its drawbacks and present a novel scheme called AHG to tackle them. Finally, we demonstrate the performance superiority of AHG over our competitor, and its viability in applications demanding the highest level of privacy.

Industry Session 2: Advances in Database Systems

Wednesday, 11:00 – 12:30

FlashStore: High Throughput Persistent Key-Value Store

Biplob Debnath (University of Minnesota, USA), Sudipta Sengupta (Microsoft Research, Redmond, USA), Jin Li (Microsoft Research, Redmond, USA)

We present FlashStore, a high throughput persistent key-value store, that uses flash memory as a non-volatile cache between RAM and hard disk. FlashStore is designed to store the working set of key-value pairs on flash and use one flash read per key lookup. As the working set changes over time, space is made for the current working set by destaging recently unused key-value pairs to hard disk and recycling pages in the flash store. FlashStore organizes key-value pairs in a log-structure on flash to exploit faster sequential write performance. It

uses an in-memory hash table to index them, with hash collisions resolved by a variant of cuckoo hashing. The in-memory hash table stores compact key signatures instead of full keys so as to strike tradeoffs between RAM usage and false flash read operations.
FlashStore can be used as a high throughput persistent key-value storage layer for a broad range of server class applications. We compare FlashStore with BerkeleyDB, an embedded key-value store application, running on hard disk and flash separately, so as to bring out the performance gain of FlashStore in not only using flash as a cache above hard disk but also in its use of flash aware algorithms. We use real-world data traces from two data center applications, namely, Xbox LIVE Primetime online multi-player game and inline storage deduplication, to drive and evaluate the design of FlashStore on traditional and low power server platforms. FlashStore outperforms BerkeleyDB by up to 60x on throughput (ops/sec), up to 50x on energy efficiency (ops/Joule), and up to 85x on cost efficiency (ops/sec/dollar) on the evaluated datasets.

MEET DB2: Automated Database Migration Evaluation

Reynold Xin (University of California Berkeley, USA), Patrick Dantressangle (IBM Hursley, United Kingdom), Sam Lightstone (IBM Research, Canada), William McLaren (IBM, USA), Steve Schormann (IBM Research, Canada), Maria Schwenger (IBM, USA)

Commercial databases compete for market share, which is composed of not only net-new sales to those purchasing a database for the first time, but also competitive "win-backs" and migrations. Database migration, or the act of moving both application code and its underlying database platform from one database to another, presents a serious administrative and application development challenge fraught with large manual costs. Migration is typically a high cost effort due to incompatibilities between database platforms. Incompatibilities are caused most often by product specific extensions to language support, procedural logic, DDL, and administrative interfaces. The migration evaluation is the first step in any competitive database migration process. Historically this has been a manual process, with the high costs and subjective results. This has led us to re-examine traditional practices and explore an automatic, innovative solution.
We have designed and implemented the Migration Evaluation and Enablement Tool for DB2 for Linux Unix and Windows, or MEET DB2, a tool for automatically evaluating database migration projects. Encapsulated in a simple one-click interface, MEET DB2 is able to provide detailed evaluation of migration complexity based on its deep analysis on the source database. In this paper, we present MEET DB2, and discuss many aspects of our design, and report measurements from real-world use cases. In particular, we show a novel way to use XML and XQuery in this domain for better extensibility and interoperability. We have evaluated MEET DB2 on 18 source code samples, covering nearly 1 million lines of code. The utility has provided benefits in several dimensions including: dramatically reduced time for evaluation, consistency, improved accuracy over human analysis, improved reporting, reduced skill requirements for migration analysis, and clear analytics for product planning.

SSD Bufferpool Extensions for Database Systems

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

High-end solid state disks (SSDs) provide much faster access to data compared to conventional hard disk drives. We present a technique for using solid-state storage as a caching layer between RAM and hard disks in database management systems. By caching data that is accessed frequently, disk I/O is reduced. For random I/O, the potential performance gains are particularly significant. Our system continuously monitors the disk access patterns to identify hot regions of the disk. Temperature statistics are maintained at the granularity of an extent, i.e., 32 pages, and are kept current through an aging mechanism. Unlike prior caching methods, once the SSD is populated with pages from warm regions cold pages are not admitted into the cache, leading to low levels of cache pollution. Simulations based on DB2 I/O traces, and a prototype implementation within DB2 both show substantial performance improvements.

Research Session 19: Probabilistic and Uncertain Databases
Wednesday, 14:00 – 15:30

UPI: A Primary Index for Uncertain Databases
Hideaki Kimura (Brown University, USA), Samuel Madden (Massachusetts Institute of Technology, USA), Stan Zdonik (Brown University, USA)

Uncertain data management has received growing attention from industry and academia. Many efforts have been made to optimize uncertain databases, including the development of special index data structures. However, none of these efforts have explored primary (clustered) indexes for uncertain databases, despite the fact that clustering has the potential to offer substantial speedups for non-selective analytic queries on large uncertain databases. In this paper, we propose a new index called a UPI (Uncertain Primary Index) that clusters heap files according to uncertain attributes with both discrete and continuous uncertainty distributions.
Because uncertain attributes may have several possible values, a UPI on an uncertain attribute duplicates tuple data once for each possible value. To prevent the size of the UPI from becoming unmanageable, its size is kept small by placing low-probability tuples in a special Cutoff Index that is consulted only when queries for low-probability values are run. We also propose several other optimizations, including techniques to improve secondary index performance and techniques to reduce maintenance costs and fragmentation by buffering changes to the table and writing updates in sequential batches. Finally, we develop cost models for UPIs to estimate query performance in various settings to help automatically select tuning parameters of a UPI.
We have implemented a prototype UPI and experimented on two real datasets. Our results show that UPIs can significantly (up to two orders of magnitude) improve the performance of uncertain queries both over clustered and unclustered attributes. We also show that our buffering techniques mitigate table fragmentation and keep the maintenance cost as low as or even lower than using an unclustered heap file.

Ranking Continuous Probabilistic Datasets
Jian Li (University of Maryland College Park, USA), Amol Deshpande (University of Maryland College Park, USA)

Ranking is a fundamental operation in data analysis and decision support, and plays an even more crucial role if the dataset being explored exhibits uncertainty. This has led to much work in understanding how to rank uncertain datasets in recent years. In this paper, we address the problem of ranking when the tuple scores are uncertain, and the uncertainty is captured using continuous probability distributions (e.g. Gaussian distributions). We present a comprehensive solution to compute the values of a parameterized ranking function (PRF) [18] for arbitrary continuous probability distributions (and thus rank the uncertain dataset); PRF can be used to simulate or approximate many other ranking functions proposed in prior work. We develop exact polynomial time algorithms for some continuous probability distribution classes, and efficient approximation schemes with provable guarantees for arbitrary probability distributions. Our algorithms can also be used for exact or approximate evaluation of k-nearest neighbor queries over uncertain objects, whose positions are modeled using continuous probability distributions. Our experimental evaluation over several datasets illustrates the effectiveness of our approach at efficiently ranking uncertain datasets with continuous attribute uncertainty.

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)

Set similarity join has played an important role in many real-world applications such as data cleaning, near duplication detection, data integration, and so on. In these applications, set data often contain noises and are thus uncertain and imprecise. In this paper, we model such probabilistic set data on two uncertainty levels, that is, set and element levels. Based on them, we investigate the problem of *probabilistic set similarity join* (PS^2J) over two probabilistic set databases, under the *possible worlds* semantics. To efficiently process the PS^2J operator, we first reduce our problem by condensing the possible worlds, and then propose effective pruning techniques, including Jaccard distance pruning, probability upper bound pruning, and aggregate pruning, which can filter out false alarms of probabilistic set pairs, with the help of indexes and our designed synopses. We demonstrate through extensive experiments the PS^2J processing performance on both real and synthetic data.

Research Session 20: Databases on Modern Hardware
Wednesday, 14:00 – 15:30

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)

Complex event detection is an advanced form of data stream processing where the stream(s) are scrutinized to identify given event patterns. The challenge for many complex event processing (CEP) systems is to be able to evaluate event patterns on high-volume data streams while adhering to real-time constraints. To solve this problem, in this paper we present a hardware-based complex event detection system implemented on field-programmable gate arrays (FPGAs). By inserting the FPGA directly into the data path between the network interface and the CPU, our solution can detect complex events at gigabit wire speed with constant and fully predictable latency, independently of network load, packet size, or data distribution. This is a significant improvement over CPU-based systems and an architectural approach that opens up interesting opportunities for hybrid stream engines that combine the exibility of the CPU with the parallelism and processing power of FPGAs.

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)

Query co-processing on graphics processors (GPUs) has become an effective means to improve the performance of main memory databases. However, this co-processing requires the data transfer between the main memory and the GPU memory via a low-bandwidth PCI-E bus. The overhead of such data transfer becomes an important factor, even a bottleneck, for query co-processing performance on the GPU. In this paper, we propose to use compression to alleviate this performance problem. Specifically, we implement nine lightweight compression schemes on the GPU and further study the combinations of these schemes for a better compression ratio. We design a compression planner to find the optimal combination. Our experiments demonstrate that the GPU-based compression and decompression achieved a processing speed up to 45 and 56 GB/s respectively. Using partial decompression, we were able to significantly improve GPU-based query co-processing performance. As a side product, we have integrated our GPU-based compression into MonetDB, an open source column-oriented DBMS, and demonstrated the feasibility of offloading compression and decompression to the GPU.

Aether: A Scalable Approach to Logging

Ryan Johnson (Carnegie Mellon University, USA), Ippokratis Pandis (Carnegie Mellon University, USA), Radu Stoica (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Manos Athanassoulis (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland)

The shift to multi-core hardware brings new challenges to database systems, as the software parallelism determines performance. Even though database systems traditionally accommodate simultaneous requests, a multitude of synchronization barriers serialize execution. Write-ahead logging is a fundamental, omnipresent component in ARIES-style concurrency and recovery, and one of the most important yet-to-be addressed potential bottlenecks, especially in OLTP workloads making frequent small changes to data.
In this paper, we identify four logging-related impediments to database system scalability. Each issue challenges different level in the software architecture: (a) the high volume of small-sized I/O requests may saturate the disk, (b) transactions hold locks while waiting for the log flush, (c) extensive context switching overwhelms the OS scheduler with threads executing log I/Os, and (d) contention appears as transactions serialize accesses to in-memory log data structures. We demonstrate these problems and address them with techniques that, when combined, comprise a holistic, scalable approach to logging. Our solution achieves a 20%-69% speedup over a modern database system when running log-intensive workloads, such as the TPC-B and TATP benchmarks. Moreover, it achieves log insert throughput over 1.8GB/s for small log records on a

single socket server, an order of magnitude higher than the traditional way of accessing the log using a single mutex.

Research Session 21: Data Mining
Wednesday, 14:00 – 15:30

Scalable Discovery of Best Clusters on Large Graphs
Kathy Macropol (University of California Santa Barbara, USA), Ambuj Singh (University of California Santa Barbara, USA)

The identification of clusters, well-connected components in a graph, is useful in many applications from biological function prediction to social community detection. However, finding these clusters can be difficult as graph sizes increase. Most current graph clustering algorithms scale poorly in terms of time or memory. An important insight is that many clustering applications need only the subset of best clusters, and not all clusters in the entire graph. In this paper we propose a new technique, Top Graph Clusters (TopGC), which probabilistically searches large, edge weighted, directed graphs for their best clusters in linear time. The algorithm is inherently parallelizable, and is able to find variable size, overlapping clusters. To increase scalability, a parameter is introduced that controls memory use. When compared with three other state-of-the art clustering techniques, TopGC achieves running time speedups of up to 70% on large scale real world datasets. In addition, the clusters returned by TopGC are consistently found to be better both in calculated score and when compared on real world benchmarks.

An Architecture for Parallel Topic Models
Alexander Smola (Yahoo! Research Labs, USA), Shravan Narayanamurthy (Yahoo! Research Labs, India)

This paper describes a high performance sampling architecture for inference of latent topic models on a cluster of workstations. Our system is faster than previous work by over an order of magnitude and it is capable of dealing with hundreds of millions of documents and thousands of topics.
The algorithm relies on a novel communication structure, namely the use of a distributed (key, value) storage for synchronizing the sampler state between computers. Our architecture entirely obviates the need for separate computation and synchronization phases. Instead, disk, CPU, and network are used simultaneously to achieve high performance. We show that this architecture is entirely general and that it can be extended easily to more sophisticated latent variable models such as n-grams and hierarchies.

Keyword++: A Framework to Improve Keyword Search Over Entity Databases
Dong Xin (Microsoft Research, USA), Yeye He (University of Wisconsin-Madison, USA), Venkatesh Ganti (Google, USA)

Keyword search over entity databases (*e.g.*, product, movie databases) is an important problem. Current techniques for keyword search on databases may often return *incomplete* and *imprecise* results. On the one hand, they either require that relevant entities contain all (or most) of the query keywords, or that relevant entities and the query keywords occur together in several documents from a known collection. Neither of these requirements may be satisfied for a number of user queries. Hence results for such queries are likely to be incomplete in that highly relevant entities may not be returned. On the other hand, although some returned entities contain all (or most) of the query keywords, the intention of the keywords in the query could be different from that in the entities. Therefore, the results could also be imprecise.
To remedy this problem, in this paper, we propose a general framework that can improve an existing search interface by translating a keyword query to a structured query. Specifically, we leverage the keyword to attribute value associations discovered in the results returned by the original search interface. We show empirically that the translated structured queries alleviate the above problems.

Research Session 22: Moving Object Databases
Wednesday, 14:00 – 15:30

Swarm: Mining Relaxed Temporal Moving Object Clusters
Zhenhui Li (University of Illinois at Urbana-Champaign, USA), Bolin Ding (University of Illinois at Urbana-Champaign, USA), Jiawei Han (University of Illinois at Urbana-Champaign, USA), Roland Kays (New York State Museum, USA)

Recent improvements in positioning technology make massive moving object data widely available. One important analysis is to find the moving objects that travel together. Existing methods put a strong constraint in defining moving object cluster, that they require the moving objects to stick together for *consecutive* timestamps. Our key observation is that the moving objects in a cluster may actually diverge temporarily and congregate at certain timestamps.
Motivated by this, we propose the concept of *swarm* which captures themoving objects that move within arbitrary shape of clusters for certain timestamps that are possibly nonconsecutive. The goal of our paper is to find all discriminative swarms, namely *closed swarm*. While the search space for closed swarms is prohibitively huge, we design a method, ObjectGrowth, to efficiently retrieve the answer. In ObjectGrowth, two effective pruning strategies are proposed to greatly reduce the search space and a novel closure checking rule is developed to report closed swarms on-thefly. Empirical studies on the real data as well as large synthetic data demonstrate the effectiveness and efficiency of our methods.

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)

In the last decade, spatio-temporal database research focuses on the design of effective and efficient indexing structures in support of location-based queries such as predictive range queries and nearest neighbor queries. While a variety of indexing techniques have been proposed to accelerate the processing of updates and queries, not much attention has been paid to the updating protocol, which is another important factor affecting system performance. In this paper, we propose a generic and adaptive updating protocol for moving object databases with less number of updating messages between the objects and database server, thereby reducing the overall workload of the system. In contrast to the approach adopted by most conventional moving object database systems where the exact locations and velocities last disclosed are used to predict their motions, we propose the concept of *Spatio-Temporal Safe Region* to approximate possible future locations. Spatio-temporal safe regions provide larger space of tolerance for moving objects, freeing them from location and velocity updates as long as the errors remain predictable in the database. To answer predictive queries accurately, the server is allowed to probe the latest status of some moving objects when their safe regions are inadequate in returning the exact query results. Spatio-temporal safe regions are calculated and optimized by the database server with two contradictory objectives: reducing update workload while guaranteeing query accuracy and efficiency. To achieve this, we propose a cost model that estimates the composition of active and passive updates based on historical motion records and query distribution. We have conducted extensive experiments to evaluate our proposal on a variety of popular indexing structures. The results confirm the viability, robustness, accuracy and efficiency of our proposed protocol.

Shortest Path Computation on Air Indexes
Georgios Kellaris (Singapore Management University, Republic of Singapore), Kyriakos Mouratidis (Singapore Management University, Republic of Singapore)

Shortest path computation is one of the most common queries in location-based services that involve transportation networks. Motivated by scalability challenges faced in the mobile network industry, we propose adopting the wireless broadcast model for such location-dependent applications. In this model the data are continuously transmitted on the air, while clients listen to the broadcast and process their queries locally. Although spatial problems have been considered in this environment, there exists no study on shortest path queries in road networks. We develop the first framework to compute shortest paths on the air, and demonstrate the practicality and efficiency of our techniques through experiments with real road networks and actual device specifications.

Research Session 23: Probabilistic Data
Wednesday, 16:00 – 17:00

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)

Probabilistic data is coming as a new deluge along with the technical advances on geographical tracking, multimedia processing, sensor network and RFID. While similarity search is an important functionality supporting the manipulation of probabilistic data, it raises new challenges to traditional relational database. The problem stems from the limited effectiveness of the distance metric supported by the existing database system. On the other hand, some complicated distance operators have proven their values for better distinguishing ability in the probabilistic domain. In this paper, we discuss the similarity search problem with the Earth Mover's Distance, which is the most successful distance metric on probabilistic histograms and an expensive operator with cubic complexity. We present a new database approach to answer range queries and k-nearest neighbour queries on probabilistic data, on the basis of Earth Mover's Distance. Our solution utilizes the primal-dual theory in linear programming and deploys B_+ tree index structures for effective candidate pruning. Extensive experiments show that our proposal dramatically improves the scalability of probabilistic databases.

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)

We show how Recursive Markov Chains (RMCs) and their restrictions can define probabilistic distributions over XML documents, and study tractability of querying over such models. We show that RMCs subsume several existing probabilistic XML models. In contrast to the latter, RMC models (i) capture probabilistic versions of XML schema languages such as DTDs, (ii) can be exponentially more succinct, and (iii) do not restrict the domain of probability distributions to be finite. We investigate RMC models for which tractability can be achieved, and identify several tractable fragments that subsume known tractable probabilistic XML models. We then look at the space of models between existing probabilistic XML formalisms and RMCs, giving results on the expressiveness and succinctness of RMC subclasses, both with each other and with prior formalisms.

Research Session 24: Fuzzy, Probabilistic and Approximate Databases
Wednesday, 16:00 – 17:00

MCDB-R: Risk Analysis in the Database

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

Enterprises often need to assess and manage the risk arising from uncertainty in their data. Such uncertainty is typically modeled as a probability distribution over the uncertain data values, specified by means of a complex (often predictive) stochastic model. The probability distribution over data values leads to a probability distribution over database query results, and risk assessment amounts to exploration of the upper or lower tail of a query-result distribution. In this paper, we extend the Monte Carlo Database System to efficiently obtain a set of samples from the tail of a query-result distribution by adapting recent "Gibbs cloning" ideas from the simulation literature to a database setting.

Scalable Probabilistic Databases with Factor Graphs and MCMC

Michael Wick (University of Massachusetts Amherst, USA), Andrew McCallum (University of Massachusetts Amherst, USA), Gerome Miklau (University of Massachusetts Amherst, USA)

Incorporating probabilities into the semantics of incomplete databases has posed many challenges, forcing systems to sacrifice modelling power, scalability, or treatment of relational algebra operators. We propose an alternative approach where the underlying relational database always represents a single world, and an external factor graph encodes a distribution over possible worlds; Markov chain Monte Carlo (MCMC) inference is then used to recover this uncertainty to a desired level of fidelity. Our approach allows the efficient evaluation of arbitrary queries over probabilistic databases with arbitrary dependencies expressed by graphical models with structure that changes during inference. MCMC sampling provides efficiency by hypothesizing modifications to possible worlds rather than generating entire worlds from scratch. Queries are then run over the portions of the world that change, avoiding the onerous cost of running full queries over each sampled world. A significant innovation of this work is the connection between MCMC sampling and materialized view maintenance techniques: we find empirically that using view maintenance techniques is several orders of magnitude faster than naively querying each sampled world. We also demonstrate our system's ability to answer relational queries with aggregation, and demonstrate additional scalability through the use of parallelization on a real-world complex model of information extraction. This framework is sufficiently expressive to support probabilistic inference not only for answering queries, but also for inferring missing database content from raw evidence.

Research Session 25: Discovery and Exploration
Wednesday, 16:00 – 17:00

On Multi-Column Foreign Key Discovery

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

A foreign/primary key relationship between relational tables is one of the most important constraints in a database. From a data analysis perspective, discovering foreign keys is a crucial step in understanding and working with the data. Nevertheless, more often than not, foreign key constraints are not specified in the data, for various reasons; e.g., some associations are not known to designers but are inherent in the data, while others become invalid due to data inconsistencies. This work proposes a robust algorithm for discovering single-column and multi-column foreign keys. Previous work concentrated mostly on discovering single-column foreign keys using a variety of rules, like inclusion dependencies, column names, and minimum/maximum values. We first propose a general rule, termed Randomness, that subsumes a variety of other rules. We then develop efficient approximation algorithms for evaluating randomness, using only two passes over the data. Finally, we validate our approach via extensive experiments using real and synthetic datasets.

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)

Data ambiguity is inherent in applications such as data integration, location-based services, and sensor monitoring. In many situations, it is possible to "clean", or remove, ambiguities from these databases. For example, the GPS location of a user is inexact due to measurement errors, but context information (e.g., what a user is doing) can be used to reduce the imprecision of the location value. In order to obtain a database with a higher quality, we study how to disambiguate a database by appropriately selecting candidates to clean. This problem is challenging because cleaning involves a cost, is limited by a budget, may fail, and may not remove all ambiguities. Moreover, the statistical information about how likely database objects can be cleaned may not be precisely known. We tackle these challenges by proposing two types of algorithms. The first type makes use of greedy heuristics to make sensible decisions; however, these algorithms do not make use of cleaning information and require user input for parameters to achieve high cleaning effectiveness. We propose the

Explore-Exploit (or *EE*) algorithm, which gathers valuable information during the cleaning process to determine how the remaining cleaning budget should be invested.
We also study how to fine-tune the parameters of *EE* in order to achieve optimal cleaning effectiveness. Experimental evaluations on real and synthetic datasets validate the effectiveness and efficiency of our approaches.

Research Session 26: Information Filtering and Dissemination
Wednesday, 16:00 – 17:00

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)

Mashups are situational applications that join multiple sources to better meet the information needs of Web users. Web sources can be huge databases behind query interfaces, which triggers the need of ranking mashup results based on some user preferences.
We present MashRank, a mashup authoring and processing system building on concepts from rank-aware processing, probabilistic databases, and information extraction to enable ranked mashups of (unstructured) sources with uncertain ranking attributes. MashRank is based on new semantics, formulations and processing techniques to handle uncertain preference scores, represented as intervals enclosing possible score values.
MashRank integrates information extraction with query processing by asynchronously pushing extracted data on-the-fly into pipelined rank-aware query plans, and using ranking early-out requirements to limit extraction cost. To the best of our knowledge, both the technical problems and target applications of MashRank have not been addressed before.

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

In this paper, we tackle the problem of efficient skycube computation. We introduce a novel approach significantly reducing domination tests for a given subspace and the number of subspaces searched. Technically, we identify two types of skyline points that can be directly derived without using any domination tests. Moreover, based on formal concept analysis, we introduce two closure operators that enable a concise representation of skyline cubes. We show that this concise representation is easy to compute and develop an efficient algorithm, which only needs to search a small portion of the huge search space. We show with empirical results the merits of our approach.

Research Session 27: Query Processing II
Thursday, 11:00 – 12:30

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)

To evaluate the performance of database applications and DBMSs, we usually execute workloads of queries on generated databases of different sizes and measure the response time. This paper introduces MyBenchmark, an offline data generation tool that takes a set of queries as input and generates database instances for which the users can control the characteristics of the resulting workload. Applications of MyBenchmark include database testing, database application testing, and application-driven benchmarking. We present the architecture and the implementation algorithms of MyBenchmark. We also present the evaluation results of MyBenchmark using TPC workloads.

Processing Top-k Join Queries

Minji Wu (Rutgers University, USA), Laure Berti-Equille (University of Rennes 1, France), Amélie Marian (Rutgers University, USA), Cecilia Procopiuc (AT&T Labs - Research, USA), Divesh Srivastava (AT&T Labs - Research, USA)

We consider the problem of efficiently finding the top-k answers for join queries over web-accessible databases. Classical algorithms for finding top-k answers use branch-and-bound techniques to avoid computing scores of all candidates in identifying the top-k answers. To be able to apply such techniques, it is critical to efficiently compute (lower and upper) bounds and expected scores of candidate answers in an incremental fashion during the evaluation.
In this paper, we describe novel techniques for these problems. The first contribution of this paper is a method to efficiently compute bounds for the score of a query result when tuples in tables from the "FROM" clause are discovered incrementally, through either sorted or random access. Our second contribution is an algorithm that, given a set of partially evaluated candidate answers, determines a good order in which to access the tables to minimize wasted efforts in the computation of top-k answers. We evaluate our algorithms on a variety of queries and data sets and demonstrate the significant benefits they provide.

Two-way Replacement Selection

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

The performance of external sorting using merge sort is highly dependent on the length of the runs generated. One of the most commonly used run generation strategies is Replacement Selection (RS) because, on average, it generates runs that are twice the size of the memory available. However, the length of the runs generated by RS is downsized for data with certain characteristics, like inputs sorted inversely with respect to the desired output order.
The goal of this paper is to propose and analyze two-way replacement selection (2WRS), which is a generalization of RS obtained by implementing two heaps instead of the single heap implemented by RS. The appropriate management of these two heaps allows generating runs larger than the memory available in a stable way, i.e. independent from the characteristics of the datasets. Depending on the changing characteristics of the input dataset, 2WRS assigns a new data record to one or the other heap, and grows or shrinks each heap, accommodating to the growing or decreasing tendency of the dataset. On average, 2WRS creates runs of at least the length generated by RS, and longer for datasets that combine increasing and decreasing data subsets. We tested both algorithms on large datasets with different characteristics and 2WRS achieves speedups at least similar to RS, and over 2.5 when RS fails to generate large runs.

Research Session 28: XML Data
Thursday, 11:00 – 12:30

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)

Previous work reports about SXSI, a fast XPath engine which executes tree automata over compressed XML indexes. Here, reasons are investigated why SXSI is so fast. It is shown that tree automata can be used as a general framework for fine grained XML query optimization. We define the "relevant nodes" of a query as those nodes that a minimal automaton must touch in order to answer the query. This notion allows to skip many subtrees during execution, and, with the help of particular tree indexes, even allows to skip internal nodes of the tree. We efficiently approximate runs over relevant nodes by means of on-the-fly removal of alternation and non-determinism of (alternating) tree automata. We also introduce many implementation techniques which allows us to efficiently evaluate tree automata, even in the absence of special indexes. Through extensive experiments, we demonstrate the impact of the different optimization techniques.

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)

In XML search systems twig queries specify predicates on node values and on the structural relationships between nodes, and a key operation is to join individual query node matches into full twig matches. Linear time twig join algorithms exist, but many non-optimal algorithms with better average-case performance have been introduced recently. These use somewhat simpler data structures that are faster in practice, but have exponential worst-case time complexity. In this paper we explore and extend the solution space spanned by previous approaches. We introduce new data structures and improved strategies for filtering out useless data nodes, yielding combinations that are both worst-case optimal and faster in practice. An experimental study shows that our best algorithm outperforms previous approaches by an average factor of three on common benchmarks. On queries with at least one unselective leaf node, our algorithm can be an order of magnitude faster, and it is never more than 20% slower on any tested benchmark query.

Destabilizers and Independence of XML Updates

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

Independence analysis is the problem of determining whether an update affects the result of a query, e.g. a constraint or materialized view. We develop a new, modular framework for static independence analysis that decomposes the problem into two orthogonal subproblems: approximating the destabilizer, that is, a finite representation of the set of updates that can change the result of the query, and testing whether the update and destabilizer overlap via an intersection analysis. Focusing on XML queries as the view language and the XQuery Update Facility as the update language, we present a syntactic query rewriting algorithm for translating queries to destabilizers, and show that intersection checking can be reduced to satisfiability problems for which efficient checkers already exist. We present an implementation based on an expressive tree satisfiability checker and a Satisfiability Modulo Order package, and give experiments confirming that the resulting analysis is both fast and effective.

Research Session 29: Workflows, Transactions, and Business Processes

Thursday, 11:00 – 12:30

Searching Workflows with Hierarchical Views

Ziyang Liu (Arizona State University, USA), Qihong Shao (Arizona State University, USA), Yi Chen (Arizona State University, USA)

Workflows are prevalent in diverse applications, which can be scientific experiments, business processes, web services, or recipes. With the dramatically growing number of workflows, there is an increasing need for people to search a workflow repository using keywords and to retrieve the relevant ones. A workflow hierarchy is a three dimensional object containing multiple abstraction views of different granularity on the same workflow. This unique structure poses a new set of challenges compared to keyword search on tree or graph structures typically found in relational or XML data.

In this paper, we define an informative, self-contained and concise search result on workflows to be a projection of a workflow hierarchy on a two dimensional viewing plane inferred from user queries. We then design and develop an efficient keyword search engine for workflows. Experimental evaluation demonstrates the effectiveness of our approach.

Data-Oriented Transaction Execution

Ippokratis Pandis (Carnegie Mellon University, USA), Ryan Johnson (Carnegie Mellon University, USA), Nikos Hardavellas (Northwestern University, USA), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland)

While hardware technology has undergone major advancements over the past decade, transaction processing systems have remained largely unchanged. The number of cores on a chip grows exponentially, following Moore's Law, allowing for an ever-increasing number of transactions to execute in parallel. As the number of concurrently-executing transactions increases, contended critical sections become scalability burdens. In typical transaction processing systems the centralized lock manager is often the first contended component and scalability bottleneck.
In this paper, we identify the conventional thread-to-transaction assignment policy as the primary cause of contention. Then, we design DORA, a system that decomposes each transaction to smaller actions and assigns actions to threads based on which data each action is about to access. DORA's design allows each thread to mostly access thread-local data structures, minimizing interaction with the contention-prone centralized lock manager. Built on top of a conventional storage engine, DORA maintains all the ACID properties. Evaluation of a prototype implementation of DORA on a multicore system demonstrates that DORA attains up to 4.8x higher throughput than a state-of-the-art storage engine when running a variety of synthetic and real-world OLTP workloads.

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, USA), Tom Yam (Tel Aviv University, Israel)

A Business Process (BP for short) consists of a set of activities that achieve some business goal when combined in a flow. Among all the (maybe infinitely many) possible execution flows of a BP, analysts are often interested in identifying flows that are "most important", according to some weight metric. This paper studies the following problem: given a specification of such a BP, a weighting function over BP execution flows, a query, and a number *k*, identify the *k* flows with the highest weight among those satisfying the query. We provide here, for the first time, a *provably optimal* algorithm for identifying the top-k weighted flows of a given BP, and use it for efficient top-k query evaluation.

Research Session 30: Scientific Databases
Thursday, 11:00 – 12:30

Behavioral Simulations in MapReduce

Guozhang Wang (Cornell University, USA), Marcos Vaz Salles (Cornell University, USA), Benjamin Sowell (Cornell University, USA), Xun Wang (Cornell University, USA), Tuan Cao (Cornell University, USA), Alan Demers (Cornell University, USA), Johannes Gehrke (Cornell University, USA), Walker White (Cornell University, USA)

In many scientific domains, researchers are turning to large-scale behavioral simulations to better understand real-world phenomena. While there has been a great deal of work on simulation tools from the high-performance computing community, behavioral simulations remain challenging to program and automatically scale in parallel environments. In this paper we present BRACE (Big Red Agent-based Computation Engine), which extends the MapReduce framework to process these simulations efficiently across a cluster. We can leverage spatial locality to treat behavioral simulations as iterated spatial joins and greatly reduce the communication between nodes. In our experiments we achieve nearly linear scale-up on several realistic simulations.
Though processing behavioral simulations in parallel as iterated spatial joins can be very efficient, it can be much simpler for the domain scientists to program the behavior of a single agent. Furthermore, many simulations include a considerable amount of complex computation and message passing between agents, which makes it important to optimize the performance of a single node and the communication across nodes. To address both of these challenges, BRACE includes a high-level language called BRASIL (the Big Red Agent SImulation Language). BRASIL has object-oriented features for programming simulations, but can be compiled to a dataflow representation for automatic parallelization and optimization. We show that by using various optimization techniques, we can achieve both scalability and single-node performance similar to that of a hand-coded simulation.

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

Tingjian Ge (University of Kentucky, USA), Stan Zdonik (Brown University, USA)

Multidimensional array database systems are suited for scientific and engineering applications. Data in these applications is often uncertain and imprecise due to errors in the instruments and observations, etc. There are often correlations exhibited in the distribution of values among the cells of an array. Typically, the correlation is stronger for cells that are close to each other and weaker for cells that are far away. We devise a novel data structure, called the *A*-tree* (*multi*dimensional *A*rray *tree*), demonstrating that by taking advantage of the predictable and structured correlations of multidimensional data, we can have a more efficient way of modeling and answering queries on large-scale array data. An A*-tree is a unified model for storage and inference. The graphical model that is assumed in an A*-tree is essentially a Bayesian Network. We analyze and experimentally verify the accuracy of an A*-tree encoding of the underlying joint distribution. We also study the efficiency of query processing over A*-trees, comparing it to an alternative graphical model.

On Dense Pattern Mining in Graph Streams

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

Many massive web and communication network applications create data which can be represented as a massive sequential stream of edges. For example, conversations in a telecommunication network or messages in a social network can be represented as a massive stream of edges. Such streams are typically very large, because of the large amount of underlying activity in such networks. An important application in these domains is to determine frequently occurring dense structures in the underlying graph stream. In general, we would like to determine frequent and dense patterns in the underlying interactions. We introduce a model for dense pattern mining and propose probabilistic algorithms for determining such structural patterns effectively and efficiently. The purpose of the probabilistic approach is to create a summarization of the graph stream, which can be used for further pattern mining. We show that this summarization approach leads to effective and efficient results for stream pattern mining over a number of real and synthetic data sets.

Research Session 31: Mobility and Spatial Queries

Thursday, 11:00 – 12:30

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 Saltenis (Aalborg University, Denmark), Kostas Tzoumas (Aalborg University, Denmark)

Given a set of users, their friend relationships, and a distance threshold per friend pair, the proximity detection problem is to find each pair of friends such that the Euclidean distance between them is within the given threshold. This problem plays an essential role in friend-locator applications and massively multiplayer online games. Existing proximity detection solutions either incur substantial location update costs or their performance does not scale well to a large number of users. Motivated by this, we present a centralized proximity detection solution that assigns each mobile client with a mobile region. We then design a *self-tuning* policy to adjust the radius of the region automatically, in order to minimize communication cost. In addition, we analyze the communication cost of our solutions, and provide valuable insights on their behaviors. Extensive experiments suggest that our proposed solution is efficient and robust with respect to various parameters.

k-Nearest Neighbors in Uncertain Graphs

Michalis Potamias (Boston University, USA), Francesco Bonchi (Yahoo! Research Labs, USA), Aristides Gionis (Yahoo! Research Labs, USA), George Kollios (Boston University, USA)

Complex networks, such as biological, social, and communication networks, often entail uncertainty, and thus, can be modeled as probabilistic graphs. Similar to the problem of similarity search in standard graphs, a

fundamental problem for probabilistic graphs is to efficiently answer k-nearest neighbor queries (k-NN), which is the problem of computing the k closest nodes to some specific node.
In this paper we introduce a framework for processing k-NN queries in probabilistic graphs. We propose novel distance functions that extend well-known graph concepts, such as shortest paths. In order to compute them in probabilistic graphs, we design algorithms based on sampling. During k-NN query processing we efficiently prune the search space using novel techniques.
Our experiments indicate that our distance functions outperform previously used alternatives in identifying true neighbours in real-world biological data. We also demonstrate that our algorithms scale for graphs with tens of millions of edges.

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)

With the increasing deployment and use of GPS-enabled devices, massive amounts of GPS data are becoming available. We propose a general framework for the mining of semantically meaningful, significant locations, e.g., shopping malls and restaurants, from such data.
We present techniques capable of extracting semantic locations from GPS data. We capture the relationships between locations and between locations and users with a graph. Significance is then assigned to locations using random walks over the graph that propagates significance among the locations. In doing so, mutual reinforcement between location significance and user authority is exploited for determining significance, as are aspects such as the number of visits to a location, the durations of the visits, and the distances users travel to reach locations. Studies using up to 100 million GPS records from a confined spatio-temporal region demonstrate that the proposal is effective and is capable of outperforming baseline methods and an extension of an existing proposal.

Industry Session 3: Data Warehousing and Mining
Thursday, 11:00 – 12:30

DataGarage: Warehousing Massive Amounts of Performance Data on Commodity Servers

Charles Loboz (Microsoft Corporation, USA), Slawek Smyl (Microsoft Corporation, USA), Suman Nath (Microsoft Research, USA)

Contemporary datacenters house tens of thousands of servers. The servers are closely monitored for operating conditions and utilizations by collecting their performance data (e.g., CPU utilization). In this paper, we show that existing database and file-system solutions are not suitable for warehousing performance data collected from a large number of servers because of the scale and the complexity of performance data. We describe the design and implementation of DataGarage, a performance data warehousing system that we have developed at Microsoft. DataGarage is a hybrid solution that combines benefits of DBMSs, file-systems, and MapReduce systems to address unique challenges of warehousing performance data. We describe how DataGarage allows efficient storage and analysis of years of historical performance data collected from many tens of thousands of servers—on commodity servers. We also report DataGarage's performance with a real dataset and a 32-node, 256-core shared-nothing cluster and our experience of using DataGarage at Microsoft for the last one year.

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

Songting Chen (Turn Incorporated, USA)

Large-scale data analysis has become increasingly important for many enterprises. Recently, a new distributed computing paradigm, called MapReduce, and its open source implementation Hadoop, has been widely adopted due to its impressive scalability and flexibility to handle structured as well as unstructured data. In this paper, we describe our data warehouse system, called Cheetah, built on top of MapReduce. Cheetah is designed specifically for our online advertising application to allow various simplifications and custom optimizations. First, we take a fresh look at the data warehouse schema design. In particular, we define a virtual view on top of the common star or snowflake data warehouse schema. This virtual view abstraction not only allows us to design a SQL-like but much more succinct query language, but also makes it easier to support many advanced query

processing features. Next, we describe a stack of optimization techniques ranging from data compression and access method to multi-query optimization and exploiting materialized views. In fact, each node with commodity hardware in our cluster is able to process raw data at 1GBytes/s. Lastly, we show how to seamlessly integrate Cheetah into any ad-hoc MapReduce jobs. This allows MapReduce developers to fully leverage the power of both MapReduce and data warehouse technologies.

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 (Ohio State University, USA), Wagner Meira (Universidade Federal de Minas Gerais, Brazil), Srinivasan Parthasarathy (Ohio State University, USA)

Detecting outliers in data is an important problem with interesting applications in a myriad of domains ranging from data cleaning to financial fraud detection and from network intrusion detection to clinical diagnosis of diseases. Over the last decade of research, distance-based outlier detection algorithms have emerged as a viable, scalable, parameter-free alternative to the more traditional statistical approaches.
In this paper we assess several distance-based outlier detection approaches and evaluate them. We begin by surveying and examining the design landscape of extant approaches, while identifying key design decisions of such approaches. We then implement an outlier detection framework and conduct a factorial design experiment to understand the pros and cons of various optimizations proposed by us as well as those proposed in the literature, both independently and in conjunction with one another, on a diverse set of real-life datasets. To the best of our knowledge this is the first such study in the literature. The outcome of this study is a family of state of the art distance-based outlier detection algorithms.
Our detailed empirical study supports the following observations. The combination of optimization strategies enables significant efficiency gains. Our factorial design study highlights the important fact that no single optimization or combination of optimizations (factors) always dominates on all types of data. Our study also allows us to characterize when a certain combination of optimizations is likely to prevail and helps provide interesting and useful insights for moving forward in this domain.

Research Session 32: Data Anonymization Techniques
Thursday, 14:00 – 15:30

Boosting the Accuracy of Differentially-Private Histograms Through Consistency

Michael Hay (University of Massachusetts Amherst, USA), Vibhor Rastogi (University of Washington, USA), Gerome Miklau (University of Massachusetts Amherst, USA), Dan Suciu (University of Washington, USA)

We show that it is possible to significantly improve the accuracy of a general class of histogram queries while satisfying differential privacy. Our approach carefully chooses a set of queries to evaluate, and then exploits consistency constraints that should hold over the noisy output. In a post-processing phase, we compute the consistent input most likely to have produced the noisy output. The final output is differentially-private and consistent, but in addition, it is often much more accurate. We show, both theoretically and experimentally, that these techniques can be used for estimating the degree sequence of a graph very precisely, and for computing a histogram that can support arbitrary range queries accurately.

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 (National University of Singapore, Republic of Singapore), Kian-Lee Tan (National University of Singapore, Republic of Singapore)

The publication of transaction data, such as market basket data, medical records, and query logs, serves the public benefit. Mining such data allows for the derivation of association rules that connect certain items to others with measurable confidence. Still, this type of data analysis poses a privacy threat; an adversary having partial information on a person's behavior may confidently associate that person to an item deemed to be sensitive. Ideally, an anonymization of such data should lead to an inference-proof version that prevents the association of individuals to sensitive items, while otherwise allowing for truthful associations to be derived. Original approaches to this problem were based on value perturbation, damaging data integrity. Recently, value generalization has been proposed as an alternative; still, approaches based on it have assumed either that all

items are equally sensitive, or that some are sensitive and can be known to an adversary only by association, while others are non-sensitive and can be known directly. Yet in reality there is a distinction between sensitive and non-sensitive items, but an adversary may possess information on any of them. Most critically, no antecedent method aims at a clear inference-proof privacy guarantee. In this paper, we propose ρ-uncertainty, the first, to our knowledge, privacy concept that inherently safeguards against sensitive associations without constraining the nature of an adversary's knowledge and without falsifying data. The problem of achieving ρ-uncertainty with low information loss is challenging because it is natural. A trivial solution is to suppress all sensitive items. We develop more sophisticated schemes. In a broad experimental study, we show that the problem is solved non-trivially by a technique that combines generalization and suppression, which also achieves favourable results compared to a baseline perturbation-based scheme.

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

Graham Cormode (AT&T Labs - Research, USA), Ninghui Li (Purdue University, USA), Tiancheng Li (Purdue University, USA), Divesh Srivastava (AT&T Labs - Research, USA)

The principle of anonymization for data sharing has become a very popular paradigm for the preservation of privacy of the data subjects. Since the introduction of k-anonymity, dozens of methods and enhanced privacy definitions have been proposed. However, over-eager attempts to minimize the information lost by the anonymization potentially allow private information to be inferred. Proof-of-concept of this "minimality attack" has been demonstrated for a variety of algorithms and definitions [16].
In this paper, we provide a comprehensive analysis and study of this attack, and demonstrate that with care its effect can be almost entirely countered. The attack allows an adversary to increase his (probabilistic) belief in certain facts about individuals over the data. We show that (a) a large class of algorithms are not affected by this attack, (b) for a class of algorithms that have a "symmetric" property, the attacker's belief increases by at most a small constant, and (c) even for an algorithm chosen to be highly susceptible to the attack, the attacker's belief when using the attack increases by at most a small constant factor. We also provide a series of experiments that show in all these cases that the confidence about the sensitive value of any individual remains low in practice, while the published data is still useful for its intended purpose. From this, we conclude that the impact of such method-based attacks can be minimized.

Research Session 33: Querying and Integrating Probabilistic Databases
Thursday, 14:00 – 15:30

Querying Probabilistic Information Extraction

Daisy Zhe Wang (University of California Berkeley, USA), Michael Franklin (University of California Berkeley, USA), Minos Garofalakis (Technical University of Crete, Greece), Joseph Hellerstein (University of California Berkeley, USA)

Recently, there has been increasing interest in extending relational query processing to include data obtained from unstructured sources. A common approach is to use stand-alone Information Extraction (IE) techniques to identify and label entities within blocks of text; the resulting entities are then imported into a standard database and processed using relational queries. This two-part approach, however, suffers from two main drawbacks. First, IE is inherently probabilistic, but traditional query processing does not properly handle probabilistic data, resulting in reduced answer quality. Second, performance inefficiencies arise due to the separation of IE from query processing. In this paper, we address these two problems by building on an in-database implementation of a leading IE model—Conditional Random Fields using the Viterbi inference algorithm. We develop two different query approaches on top of this implementation. The first uses deterministic queries over maximum-likelihood extractions, with optimizations to push the relational operators into the Viterbi algorithm. The second extends the Viterbi algorithm to produce a set of possible extraction "worlds", from which we compute top-k probabilistic query answers. We describe these approaches and explore the trade-offs of efficiency and effectiveness between them using two datasets.

Read-Once Functions and Query Evaluation in Probabilistic Databases

Prithviraj Sen (Yahoo! Research Labs, USA), Amol Deshpande (University of Maryland College Park, USA), Lise Getoor (University of Maryland College Park, USA)

Probabilistic databases hold promise of being a viable means for large-scale uncertainty management, increasingly needed in a number of real world applications domains. However, query evaluation in probabilistic databases remains a computational challenge. Prior work on efficient exact query evaluation in probabilistic databases has largely concentrated on query-centric formulations (e.g., safe plans, hierarchical queries), in that, they only consider characteristics of the query and not the data in the database. It is easy to construct examples where a supposedly hard query run on an appropriate database gives rise to a tractable query evaluation problem. In this paper, we develop efficient query evaluation techniques that leverage characteristics of both the query and the data in the database. We focus on tuple-independent databases where the query evaluation problem is equivalent to computing marginal probabilities of Boolean formulas associated with the result tuples. This latter task is easy if the Boolean formulas can be factorized into a form that has every variable appearing at most once (called read-once). However, a naive approach that directly uses previously developed Boolean formula factorization algorithms is inefficient, because those algorithms require the input formulas to be in the disjunctive normal form (DNF). We instead develop novel, more efficient factorization algorithms that directly construct the read-once expression for a result tuple Boolean formula (if one exists), for a large subclass of queries (specifically, conjunctive queries without self-joins). We empirically demonstrate that (1) our proposed techniques are orders of magnitude faster than generic inference algorithms for queries where the result Boolean formulas can be factorized into read-once expressions, and (2) for the special case of hierarchical queries, they rival the efficiency of prior techniques specifically designed to handle such queries.

Foundations of Uncertain-Data Integration

Parag Agrawal (Stanford University, USA), Anish Das Sarma (Yahoo! Research Labs, USA), Jeffrey Ullman (Stanford University, USA), Jennifer Widom (Stanford University, USA)

There has been considerable past work studying data integration and uncertain data in isolation. We develop the foundations for local-as-view (LAV) data integration when the sources being integrated are uncertain. We motivate two distinct settings for uncertain-data integration. We then define containment of uncertain databases in these settings, which allows us to express uncertain sources as views over a virtual mediated uncertain database. Next, we define consistency of a set of uncertain sources and show intractability of consistency-checking. We identify an interesting special case for which consistency-checking is polynomial. Finally, the notion of certain answers from traditional LAV data integration does not generalize to the uncertain setting, so we define a corresponding notion of correct answers.

Research Session 34: Database Design
Thursday, 14:00 – 15:30

Identifying, Attributing and Describing Spatial Bursts

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

User generated content that appears on weblogs, wikis and social networks has been increasing at an unprecedented rate. The wealth of information produced by individuals from different geographical locations presents a challenging task of intelligent processing.
In this paper, we introduce a methodology to identify notable geographically focused events out of this collection of user generated information. At the heart of our proposal lie efficient algorithms that identify geographically focused information bursts, attribute them to demographic factors and identify sets of descriptive keywords. We present the results of a prototype evaluation of our algorithms on BlogScope, a large-scale social media warehousing platform. We demonstrate the scalability and practical utility of our proposal running on top of a multi-terabyte text collection.

CORADD: Correlation Aware Database Designer for Materialized Views and Indexes

Hideaki Kimura (Brown University, USA), George Huo (Google, USA), Alexander Rasin (Brown University, USA), Samuel Madden (Massachusetts Institute of Technology, USA), Stanley Zdonik (Brown University, USA)

We describe an automatic database design tool that exploits correlations between attributes when recommending materialized views (MVs) and indexes. Although there is a substantial body of related work exploring how to select an appropriate set of MVs and indexes for a given workload, none of this work has explored the effect of correlated attributes (e.g., attributes encoding related geographic information) on designs. Our tool identifies a set of MVs and secondary indexes such that correlations between the clustered attributes of the MVs and the secondary indexes are enhanced, which can dramatically improve query performance. It uses a form of Integer Linear Programming (ILP) called ILP Feedback to pick the best set of MVs and indexes for given database size constraints. We compare our tool with a state-of-the-art commercial database designer on two workloads, APB-1 and SSB (Star Schema Benchmark—similar to TPC-H). Our results show that a correlation-aware database designer can improve query performance up to 6 times within the same space budget when compared to a commercial database designer.

Regret-Minimizing Representative Databases

Danupon Nanongkai (Georgia Institute of Technology, USA) , Atish Das Sarma (Georgia Institute of Technology, USA), Ashwin Lall (Georgia Institute of Technology, USA), Richard Lipton (Georgia Institute of Technology, USA), Jun Xu (Georgia Institute of Technology, USA)

We propose the *k*-representative regret minimization query (*k*-regret) as an operation to support multi-criteria decision making. Like top-*k*, the *k*-regret query assumes that users have some utility or scoring functions; however, it never asks the users to provide such functions. Like skyline, it filters out a set of interesting points from a potentially large database based on the users' criteria; however, it never overwhelms the users by outputting too many tuples.

In particular, for any number *k* and any class of utility functions, the *k*-regret query outputs *k* tuples from the database and tries to minimize the *maximum regret ratio*. This captures how disappointed a user could be had she seen *k* representative tuples instead of the whole database. We focus on the class of linear utility functions, which is widely applicable.

The first challenge of this approach is that it is not clear if the maximum regret ratio would be small, or even bounded. We answer this question affirmatively. Theoretically, we prove that the maximum regret ratio can be bounded and this bound is independent of the database size. Moreover, our extensive experiments on real and synthetic datasets suggest that in practice the maximum regret ratio is reasonably small. Additionally, algorithms developed in this paper are practical as they run in linear time in the size of the database and the experiments show that their running time is small when they run on top of the skyline operation which means that these algorithm could be integrated into current database systems.

Research Session 35: Database Design

Thursday, 14:00 – 15:30

An Access CostAware Approach for Object Retrieval over Multiple Sources

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

Source and object selection and retrieval from large multi-source data sets are fundamental operations in many applications. In this paper, we initiate research on efficient source (e.g., database) and object selection algorithms on large multi-source data sets. Specifically, in order to acquire a specified number of satisfying objects with minimum cost over multiple databases, the query engine needs to determine the access overhead for individual data sources, the overhead of retrieving objects from each source, and possibly other statistics such as estimating the frequency of finding a satisfying object in order to determine how many objects to retrieve from each data source. We adopt a probabilistic approach to source selection utilizing a cost structure and a dynamic programming model for computing the optimal number of objects to retrieve from each data source. Such a structure can be a valuable asset where there is a monetary or time related cost associated with accessing large distributed databases. We present a thorough experimental evaluation to validate our techniques using real-world data sets.

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)

Predicate selectivity estimates are subject to considerable run-time variation relative to their compile-time estimates, often leading to poor plan choices that cause inflated response times. We present here a parametrized family of plan generation and selection algorithms that replace, whenever feasible, the optimizer's solely costconscious choice with an alternative plan that is (a) guaranteed to be near-optimal in the absence of selectivity estimation errors, and (b) likely to deliver comparatively stable performance in the presence of arbitrary errors. These algorithms have been implemented within the PostgreSQL optimizer, and their performance evaluated on a rich spectrum of TPC-H and TPC-DS-based query templates in a variety of database environments. Our experimental results indicate that it is indeed possible to identify robust plan choices that substantially curtail the adverse effects of erroneous selectivity estimates. In fact, the plan selection quality provided by our algorithms is often competitive with those obtained through apriori knowledge of the plan search and optimality spaces. The additional computational overheads incurred by the replacement approach are miniscule in comparison to the expected savings in query execution times. We also demonstrate that with appropriate parameter choices, it is feasible to directly produce anorexic plan diagrams, a potent objective in query optimizer design.

Xplus: A SQL-Tuning-Aware Query Optimizer

Herodotos Herodotou (Duke University, USA), Shivnath Babu (Duke University, USA)

The need to improve a suboptimal execution plan picked by the query optimizer for a repeatedly run SQL query arises routinely. Complex expressions, skewed or correlated data, and changing conditions can cause the optimizer to make mistakes. For example, the optimizer may pick a poor join order, overlook an important index, use a nested-loop join when a hash join would have done better, or cause an expensive, but avoidable, sort to happen. SQL tuning is also needed while tuning multi-tier services to meet service-level objectives. The difficulty of SQL tuning can be lessened considerably if users and higher-level tuning tools can tell the optimizer: "I am not satisfied with the performance of the plan p being used for the query Q that runs repeatedly. Can you generate a (δ%) better plan?" This paper designs, implements, and evaluates Xplus which, to our knowledge, is the first query optimizer to provide this feature. Xplus goes beyond the traditional plan-first-execute-next approach: Xplus runs some (sub)plans proactively, collects monitoring data from the runs, and iterates. A nontrivial challenge is in choosing a small set of plans to run. Xplus guides this process efficiently using an extensible architecture comprising SQL-tuning experts with different goals, and a policy to arbitrate among the experts. We show the effectiveness of Xplus on real-life tuning scenarios created using TPC-H queries on a PostgreSQL database.

Research Session 36: Graph and Pattern Matching
Thursday, 14:00 – 15:30

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)

In a variety of emerging applications one needs to decide whether a graph G *matches* another G_p, *i.e.,* whether G has a topological structure similar to that of G_p. The traditional notions of graph homomorphism and isomorphism often fall short of capturing the structural similarity in these applications. This paper studies revisions of these notions, providing a full treatment from complexity to algorithms. (1) We propose p-*homomorphism* (p-hom) and 1-1 p-hom, which extend graph homomorphism and subgraph isomorphism, respectively, by mapping *edges* from one graph to *paths* in another, and by measuring *the similarity of nodes*. (2) We introduce metrics to measure graph similarity, and several optimization problems for p-hom and 1-1 p-hom. (3) We show that the decision problems for p-hom and 1-1 p-hom are NP-complete even for DAGs, and that the optimization problems are approximation-hard. (4) Nevertheless, we provide approximation algorithms

with *provable guarantees* on match quality. We experimentally verify the effectiveness of the revised notions and the efficiency of our algorithms in Web site matching, using real-life and synthetic data.

SigMatch: Fast and Scalable Multi-Pattern Matching

Ramakrishnan Kandhan (University of Wisconsin-Madison, USA), Nikhil Teletia (University of Wisconsin-Madison, USA), Jignesh Patel (University of Wisconsin-Madison, USA)

Multi-pattern matching involves matching a data item against a large database of "signature" patterns. Existing algorithms for multi-pattern matching do not scale well as the size of the signature database increases. In this paper, we present sigMatch – a fast, versatile, and scalable technique for multi-pattern signature matching. At its heart, sigMatch organizes the signature database into a (processor) cache-efficient q-gram index structure, called the sigTree. The sigTree groups patterns based on common sub-patterns, such that signatures that don't match can be quickly eliminated from the matching process. The sigTree also uses parallel Bloom filters and a technique to reduce imbalances across groups, for improved performance. Using extensive empirical evaluation across three diverse domains, we show that sigMatch often outperforms existing methods by an order of magnitude or more.

SAPPER: Subgraph Indexing and Approximate Matching in Large Graphs

Shijie Zhang (Case Western Reserve University, USA), Jiong Yang (Case Western Reserve University, USA), Wei Jin (Case Western Reserve University, USA)

With the emergence of new applications, e.g., computational biology, new software engineering techniques, social networks, etc., more data is in the form of graphs. Locating occurrences of a query graph in a large database graph is an important research topic. Due to the existence of noise (e.g., missing edges) in the large database graph, we investigate the problem of approximate subgraph indexing, i.e., finding the occurrences of a query graph in a large database graph with (possible) missing edges. The SAPPER method is proposed to solve this problem. Utilizing the hybrid neighborhood unit structures in the index, SAPPER takes advantage of pre-generated random spanning trees and a carefully designed graph enumeration order. Real and synthetic data sets are employed to demonstrate the efficiency and scalability of our approximate subgraph indexing method.

Industry Session 4: Mobile and Social Applications

Thursday, 14:00 – 15:30

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)

Nowadays, due to the increased user requirements of the fast and reliable data management operation for mobile applications, major device vendors use embedded DBMS for their mobile devices such as MP3 players, mobile phones, digital cameras and PDAs. However, database logging is the major bottleneck against the fast response time. There has been a lot of work minimizing logging overhead but no single recovery method provides the best performance to a variety of database workloads. In this paper, we present a novel recovery method called *adaptive logging* which can switch the logging method from ARIES to shadow paging adaptively at a page level according to the update state of each page on run time. Also, we propose a log compaction method called *deferred logging* which removes redundant logs by deferring to create log records until the updated data page is flushed or until the transaction commits. Deferred logging is coupled with adaptive logging seamlessly so that it boosts the performance of adaptive logging by reducing the typical overhead of hybrid methods. We have implemented the proposed approaches to our embedded DBMS which was deployed to more than 10 million mobile devices and evaluated them through a real world application on a mobile device. The result shows that our approaches can reduce logging overhead significantly and consequently can improve the response time of both small update transaction and large update transaction effectively.

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

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

Mobile commerce and location based services (LBS) are some of the fastest growing IT industries in the last five years. Location update of mobile clients is a fundamental capability in mobile commerce and all types of LBS. Higher update frequency leads to higher accuracy, but incurs unacceptably high cost of location management at the location servers. We propose ROADTRACK – a road-network based, query-aware location update framework with two unique features. First, we introduce the concept of precincts to control the granularity of location update resolution for mobile clients that are not of interest to any active location query services. Second, we define query encounter points for mobile objects that are targets of active location query services, and utilize these encounter points to define the adequate location update schedule for each mobile. The ROADTRACK framework offers three unique advantages. First, encounter points as a fundamental query awareness mechanism enable us to control and differentiate location update strategies for mobile clients in the vicinity of active location queries, while meeting the needs of location query evaluation. Second, we employ system-defined precincts to manage the desired spatial resolution of location updates for different mobile clients and to control the scope of query awareness to be capitalized by a location update strategy. Third, our road-network based check-free interval optimization further enhances the effectiveness of the ROADTRACK query-aware location update scheduling algorithm. This optimization provides significant cost reduction for location update management at both mobile clients and location servers. We evaluate the ROADTRACK location update approach using a real world road-network based mobility simulator. Our experimental results demonstrate that the ROADTRACK query aware location update approach outperforms existing representative location update strategies in terms of both client energy efficiency and server processing load.

Confucius and its Intelligent Disciples: Integrating Social with Search

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

Q&A sites continue to flourish as a large number of users rely on them as useful substitutes for incomplete or missing search results. In this paper, we present our experience with developing Confucius, a Google Q&A service launched in 21 countries and four languages by the end of 2009. Confucius employs six data mining subroutines to harness synergy between web search and social networks. We present these subroutines' design goals, algorithms, and their effects on service quality. We also describe techniques for and experience with scaling the subroutines to mine massive data sets.

Research Session 37: Indexing Techniques

Thursday, 16:00 – 17:30

Tree Indexing on Solid State Drives

Yinan Li (University of Wisconsin-Madison, USA), 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)

Large flash disks, or solid state drives (SSDs), have become an attractive alternative to magnetic hard disks, due to their high random read performance, low energy consumption and other features. However, writes, especially small random writes, on flash disks are inherently much slower than reads because of the erase-before-write mechanism.

To address this asymmetry of read-write speeds in tree indexing on the flash disk, we propose FD-tree, a tree index designed with the *logarithmic method* and *fractional cascading* techniques. With the logarithmic method, an FD-tree consists of the head tree – a small B+-tree on the top, and a few levels of sorted runs of increasing sizes at the bottom. This design is write-optimized for the flash disk; in particular, an index search will potentially go through more levels or visit more nodes, but random writes are limited to a small area – the head

tree, and are subsequently transformed into sequential ones through merging into the lower runs. With the fractional cascading technique, we store pointers, called fences, in lower level runs to speed up the search. Given an FD-tree of n entries, we analytically show that it performs an update in $O(\log_B n)$ sequential I/Os and completes a search in $O(\log_B n)$ random I/Os, where B is the flash page size. We evaluate FD-tree in comparison with representative B+-tree variants under a variety of workloads on three commodity flash SSDs. Our results show that FD-tree has a similar search performance to the standard B+-tree, and a similar update performance to the write-optimized B+-tree variant. As a result, FD-tree dominates the other B+-tree index variants on the overall performance on flash disks as well as on magnetic disks.

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, USA)

A Cloud may be seen as a type of flexible computing infrastructure consisting of many compute nodes, where resizable computing capacities can be provided to different customers. To fully harness the power of the Cloud, efficient data management is needed to handle huge volumes of data and support a large number of concurrent end users. To achieve that, a scalable and high-throughput indexing scheme is generally required. Such an indexing scheme must not only incur a low maintenance cost but also support parallel search to improve scalability. In this paper, we present a novel, scalable B+-tree based indexing scheme for efficient data processing in the Cloud. Our approach can be summarized as follows. First, we build a local B+-tree index for each compute node which only indexes data residing on the node. Second, we organize the compute nodes as a structured overlay and publish a portion of the local B+-tree nodes to the overlay for efficient query processing. Finally, we propose an adaptive algorithm to select the published B+-tree nodes according to query patterns. We conduct extensive experiments on Amazon's EC2, and the results demonstrate that our indexing scheme is dynamic, efficient and scalable.

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)

A string similarity join finds similar pairs between two collections of strings. It is an essential operation in many applications, such as data integration and cleaning, and has attracted significant attention recently. In this paper, we study string similarity joins with edit-distance constraints. Existing methods usually employ a filter-and-refine framework and have the following disadvantages: (1) They are inefficient for the data sets with short strings (the average string length is no larger than 30); (2) They involve large indexes; (3) They are expensive to support dynamic update of data sets. To address these problems, we propose a novel framework called trie-join, which can generate results efficiently with small indexes. We use a trie structure to index the strings and utilize the trie structure to efficiently find the similar string pairs based on subtrie pruning. We devise efficient trie-join algorithms and pruning techniques to achieve high performance. Our method can be easily extended to support dynamic update of data sets efficiently. Experimental results show that our algorithms outperform state-of-the-art methods by an order of magnitude on three real data sets with short strings.

Research Session 38: Query Processing III
Thursday, 16:00 – 17:30

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

Mehdi Sharifzadeh (Google Santa Monica, USA), Cyrus Shahabi (University of Southern California, USA)

A very important class of spatial queries consists of nearest-neighbor (NN) query and its variations. Many studies in the past decade utilize R-trees as their underlying index structures to address NN queries efficiently. The general approach is to use R-tree in two phases. First, R-tree's hierarchical structure is used to quickly arrive to the neighbourhood of the result set. Second, the R-tree nodes intersecting with the local neighborhood (Search Region) of an initial answer are investigated to find all the members of the result set. While R-trees are

very efficient for the first phase, they usually result in the unnecessary investigation of many nodes that none or only a small subset of their including points belongs to the actual result set.
On the other hand, several recent studies showed that the Voronoi diagrams are extremely efficient in exploring an NN search region, while due to lack of an efficient access method, their arrival to this region is slow. In this paper, we propose a new index structure, termed VoR-Tree that incorporates Voronoi diagrams into R-tree, benefiting from the best of both worlds. The coarse granule rectangle nodes of R-tree enable us to get to the search region in logarithmic time while the fine granule polygons of Voronoi diagram allow us to efficiently tile or cover the region and find the result. Utilizing VoR-Tree, we propose efficient algorithms for various Nearest Neighbor queries, and show that our algorithms have better I/O complexity than their best competitors.

Efficient RkNN Retrieval with Arbitrary Non-Metric Similarity Measures

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

A RkNN query returns all objects whose nearest k neighbours contain the query object. In this paper, we consider RkNN query processing in the case where the distances between attribute values are not necessarily metric. Dissimilarities between objects could then be a monotonic aggregate of dissimilarities between their values, such aggregation functions being specified at query time. We outline real world cases that motivate RkNN processing in such scenarios. We consider the AL-Tree index and its applicability in RkNN query processing. We develop an approach that exploits the group level reasoning enabled by the AL-Tree in RkNN processing. We evaluate our approach against a Naive approach that performs sequential scans on contiguous data and an improved block-based approach that we provide. We use real-world datasets and synthetic data with varying characteristics for our experiments. This extensive empirical evaluation shows that our approach is better than existing methods in terms of computational and disk access costs, leading to significantly better response times.

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 W. Cheung (The University of Hong Kong, People's Republic of China)

Although there has been a considerable body of work on skyline evaluation in multidimensional data with totally ordered attribute domains, there are only a few methods that consider attributes with partially ordered domains. Existing work maps each partially ordered domain to a total order and then adapts algorithms for totallyordered domains to solve the problem. Nevertheless these methods either use stronger notions of dominance, which generate false positives, or require expensive dominance checks. In this paper, we propose two new methods, which do not have these drawbacks. The first method uses an appropriate mapping of a partial order to a total order, inspired by the lattice theorem and an off-the-shelf skyline algorithm. The second technique uses an appropriate storage and indexing approach, inspired by column stores, which enables efficient verification of whether a pair of objects are incompatible. We demonstrate that both our methods are up to an order of magnitude more efficient than previous work and scale well with different problem parameters, such as complexity of partial orders.

Research Session 39: Streaming and Sensor Data

Thursday, 16:00 – 18:00

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

Mingzhu Wei (Worcester Polytechnic Institute, USA), Elke Rundensteiner (Worcester Polytechnic Institute, USA), Murali Mani (University of Michigan at Flint, USA)

Because of high volumes and unpredictable arrival rates, stream processing systems are not always able to keep up with input data - resulting in buffer overflow and uncontrolled loss of data. To produce eventually complete results, load spilling, which pushes some fractions of data to disks temporarily, is commonly employed in relational stream engines. In this work, we now introduce "structure-based spilling", a spilling technique customized for XML streams by considering the partial spillage of possibly complex XML elements. Such structure-based spilling brings new challenges. When a path is spilled, multiple paths may be affected. We

analyze possible spilling effects on the query paths and how to execute the "reduced" query to produce partial results. To select the reduced query that maximizes output quality, we develop three optimization strategies, namely, OptR, OptPrune and ToX. We also examine the clean-up stage to guarantee that an entire result set is eventually generated by producing supplementary results. Our experimental study demonstrates that our proposed solutions consistently achieve higher quality results compared to the state-of-the-art techniques.

Dynamic Join Optimization in Multi-Hop Wireless Sensor Networks

Svilen Mihaylov (University of Pennsylvania, USA), Marie Jacob (University of Pennsylvania, USA), Zachary Ives (University of Pennsylvania, USA), Sudipto Guha (University of Pennsylvania, USA)

To enable smart environments and self-tuning data centers, we are developing the Aspen system for integrating physical sensor data, as well as stream data coming from machine logical state, and database or Web data from the Internet. A key component of this system is a query processor optimized for limited-bandwidth, possibly battery-powered devices with multiple hop wireless radio communications. This query processor is given a portion of a data integration query, possibly including joins among sensors, to execute. Several recent papers have developed techniques for computing joins in sensors, but these techniques are static and are only appropriate for specific join selectivity ratios. We consider the problem of dynamic join optimization for sensor networks, developing solutions that employ cost modeling, as well as adaptive learning and self-tuning heuristics to choose the best algorithm under real and variable selectivity values. We focus on in-network join computation, but our architecture extends to other approaches (and we compare against these). We develop basic techniques assuming selectivities are uniform and known in advance, and optimization can be done on a pairwise basis; we then extend the work to handle joins between multiple pairs, when selectivities are not fully known. We experimentally validate our work at scale using standard datasets.

Database-support for Continuous Prediction Queries over Streaming Data

Mert Akdere (Brown University, USA), Ugur Cetintemel (Brown University, USA), Eli Upfal (Brown University, USA)

Prediction is emerging as an essential ingredient for real-time monitoring, planning and decision support applications such as intrusion detection, e-commerce pricing and automated resource management. This paper presents a system that efficiently supports continuous prediction queries (CPQs) over streaming data using seamlessly-integrated probabilistic models. Specifically, we describe how to execute and optimize CPQs using discrete (Dynamic) Bayesian Networks as the underlying predictive model. Our primary contribution is a novel cost-based optimization framework that employsmaterialization, sharing, andmodel-specific optimization techniques to enable highly-efficient point- and range-based CPQ execution. Furthermore, we support efficient execution of top-k and threshold-based high probability queries. We characterize the behavior of our system and demonstrate significant performance gains using a prototype implementation operating on realworld network intrusion data and deployed as part of a real-time software-performance monitoring system.

Conditioning and Aggregating Uncertain Data Streams: Going Beyond Expectations

Thanh Tran (University of Massachusetts Amherst, USA), Andrew McGregor (University of Massachusetts Amherst, USA), Yanlei Diao (University of Massachusetts Amherst, USA), Liping Peng (University of Massachusetts Amherst, USA), Anna Liu (University of Massachusetts Amherst, USA)

Uncertain data streams are increasingly common in real-world deployments and monitoring applications require the evaluation of complex queries on such streams. In this paper, we consider complex queries involving conditioning (e.g., selections and group by's) and aggregation operations on uncertain data streams. To characterize the uncertainty of answers to these queries, one generally has to compute the full probability distribution of each operation used in the query. Computing distributions of aggregates given conditioned tuple distributions is a hard, unsolved problem. Our work employs a new evaluation framework that includes a general data model, approximation metrics, and approximate representations. Within this framework we design fast data-stream algorithms, both deterministic and randomized, for returning approximate distributions with bounded errors as answers to those complex queries. Our experimental results demonstrate the accuracy and efficiency of our approximation techniques and offer insights into the strengths and limitations of deterministic and randomized algorithms.

Research Session 40: Information Integration and Retrieval
Thursday, 16:00 – 17:30

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 Almaden Research Center, USA)

Though partially automated, developing schema mappings remains a complex and potentially error-prone task. In this paper, we present TRAMP (TRAnsformation Mapping Provenance), an extensive suite of tools supporting the debugging and tracing of schema mappings and transformation queries. TRAMP combines and extends data provenance with two novel notions, transformation provenance and mapping provenance, to explain the relationship between transformed data and those transformations and mappings that produced that data. In addition we provide query support for transformations, data, and all forms of provenance. We formally define transformation and mapping provenance, present an efficient implementation of both forms of provenance, and evaluate the resulting system through extensive experiments.

Entity Resolution with Evolving Rules
Steven Whang (Stanford University, USA), Hector Garcia-Molina (Stanford University, USA)

Entity resolution (ER) identifies database records that refer to the same real world entity. In practice, ER is not a one-time process, but is constantly improved as the data, schema and application are better understood. We address the problem of keeping the ER result up-to-date when the ER logic "evolves" frequently. A naïve approach that re-runs ER from scratch may not be tolerable for resolving large datasets. This paper investigates when and how we can instead exploit previous "materialized" ER results to save redundant work with evolved logic. We introduce algorithm properties that facilitate evolution, and we propose efficient rule evolution techniques for two clustering ER models: match-based clustering and distance-based clustering. Using real data sets, we illustrate the cost of materializations and the potential gains over the naive approach.

Annotating and Searching Web Tables Using Entities, Types and Relationships
Girija Limaye (IIT Bombay, India), Sunita Sarawagi (IIT Bombay, India), Soumen Chakrabarti (IIT Bombay, India)

Tables are a universal idiom to present relational data. Billions of tables on Web pages express entity references, attributes and relationships. This representation of relational world knowledge is usually considerably better than completely unstructured, free-format text. At the same time, unlike manually-created knowledge bases, relational information mined from "organic" Web tables need not be constrained by availability of precious editorial time. Unfortunately, in the absence of any formal, uniform schema imposed on Web tables, Web search cannot take advantage of these high-quality sources of relational information. In this paper we propose new machine learning techniques to annotate table cells with entities that they likely mention, table columns with types from which entities are drawn for cells in the column, and relations that pairs of table columns seek to express. We propose a new graphical model for making all these labeling decisions for each table simultaneously, rather than make separate local decisions for entities, types and relations. Experiments using the YAGO catalog, DBPedia, tables from Wikipedia, and over 25 million HTML tables from a 500 million page Web crawl uniformly show the superiority of our approach. We also evaluate the impact of better annotations on a prototype relational Web search tool. We demonstrate clear benefits of our annotations beyond indexing tables in a purely textual manner.

Research Session 41: Data Mining, Copy Detection and Data Publication
Thursday, 16:00 – 17:30

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)

Large text corpora with news, customer mail and reports, or Web 2.0 contributions offer a great potential for enhancing business-intelligence applications. We propose a framework for performing text analytics on such data in a versatile, efficient, and scalable manner. While much of the prior literature has emphasized mining keywords or tags in blogs or social-tagging communities, we emphasize the analysis of interesting phrases. These include named entities, important quotations, market slogans, and other multi-word phrases that are prominent in a dynamically derived ad-hoc subset of the corpus, e.g., being frequent in the subset but relatively infrequent in the overall corpus. We develop preprocessing and indexing methods for phrases, paired with new search techniques for the top-k most interesting phrases in ad-hoc subsets of the corpus. Our framework is evaluated using a large-scale real-world corpus of New York Times news articles.

Global Detection of Complex Copying Relationships Between Sources

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

Web technologies have enabled data sharing between sources but also simplified copying (and often publishing without proper attribution). The copying relationships can be complex: some sources copy from multiple sources on different subsets of data; some co-copy from the same source, and some transitively copy from another. Understanding such copying relationships is desirable both for business purposes and for improving many key components in data integration, such as resolving conflicts across various sources, reconciling distinct references to the same real-world entity, and efficiently answering queries over multiple sources. Recent works have studied how to detect copying between a pair of sources, but the techniques can fall short in the presence of complex copying relationships. In this paper we describe techniques that discover global copying relationships between a set of structured sources. Towards this goal we make two contributions. First, we propose a global detection algorithm that identifies co-copying and transitive copying, returning only source pairs with direct copying. Second, global detection requires accurate decisions on copying direction; we significantly improve over previous techniques on this by considering various types of evidence for copying and correlation of copying on different data items. Experimental results on real-world data and synthetic data show high effectiveness and efficiency of our techniques.

Fragments and Loose Associations: Respecting Privacy in Data Publishing

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

We propose a modeling of the problem of privacy-compliant data publishing that captures confidentiality constraints on one side and visibility requirements on the other side. Confidentiality constraints express the fact that some attributes, or associations among them, are sensitive and cannot be released. Visibility requirements express requests for views over data that should be provided. We propose a solution based on data fragmentation to split sensitive associations while ensuring visibility. In addition, we show how sensitive associations broken by fragmentation can be released in a sanitized form as loose associations formed in a way to guarantee a specified degree of privacy.

CONFERENCE CO-LOCATED WORKSHOPS

Ph.D. Workshop
http://www.vldb2010.org/cfp-phd-workshop.htm
September 13, 2010
Location:

The VLDB 2010 Ph.D. Workshop will provide a high-quality international forum for Ph.D. students working in the areas related to VLDB conference series to present their ongoing research, share ideas, receive constructive feedback and establish professional networks with other PhD students and more senior researchers at the VLDB conference. A group of prominent researchers and high-performing early career researchers will be invited to participate the discussions at the workshop.

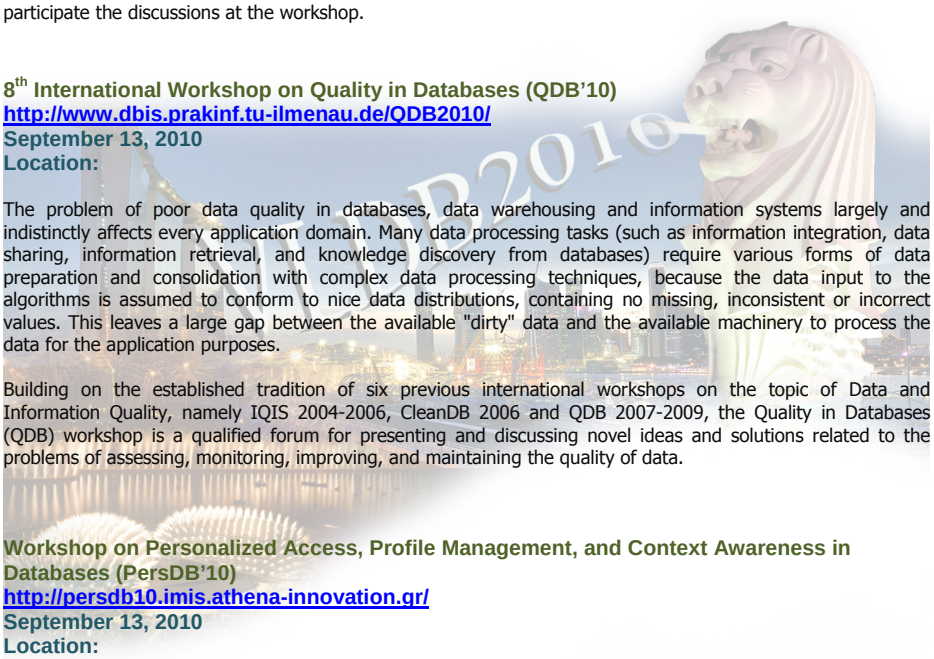

8th International Workshop on Quality in Databases (QDB'10)
http://www.dbis.prakinf.tu-ilmenau.de/QDB2010/
September 13, 2010
Location:

The problem of poor data quality in databases, data warehousing and information systems largely and indistinctly affects every application domain. Many data processing tasks (such as information integration, data sharing, information retrieval, and knowledge discovery from databases) require various forms of data preparation and consolidation with complex data processing techniques, because the data input to the algorithms is assumed to conform to nice data distributions, containing no missing, inconsistent or incorrect values. This leaves a large gap between the available "dirty" data and the available machinery to process the data for the application purposes.

Building on the established tradition of six previous international workshops on the topic of Data and Information Quality, namely IQIS 2004-2006, CleanDB 2006 and QDB 2007-2009, the Quality in Databases (QDB) workshop is a qualified forum for presenting and discussing novel ideas and solutions related to the problems of assessing, monitoring, improving, and maintaining the quality of data.

Workshop on Personalized Access, Profile Management, and Context Awareness in Databases (PersDB'10)
http://persdb10.imis.athena-innovation.gr/
September 13, 2010
Location:

Proliferation of database-driven web sites has brought upon a plethora of information access and dissemination applications. Monitoring and trading stock portfolios, news notification, weather tracking, and even simple search are just a few examples. In addition, recently, new applications have emerged that go beyond information access and dissemination and enhance creativity, information sharing, and collaboration among users providing richer interaction possibilities. Now, users cannot only access content but they can also generate, share and modify content (both theirs and others') freely, compose their applications, enhance their interface, etc. Web-based communities, wikis, social networks, mashups, folksonomies are some of the emerging new applications. In all these (classical and novel) applications, different notions of user information, such as preferences, community memberships and social interactions, and context information, such as a user's social network, location, time, and other features of a user's environment, are of paramount importance in order to improve and personalize user experience. In this context, new challenges emerge for **user-centric, context-aware database systems** for storing and managing different aspects of user and context information, for data management and computing taking into account personal, social and contextual information about users and for customizability of their behavior.

User-centric, personalized, socially-affected, and context-aware database systems represent a remarkable step towards user-centric applications that allow users to find, generate, share and modify content and services. It is

imperative that people studying and working on the different components of a database system clarify their view of contextualization and personalization and describe the ways and degree to which these can be applied to their components of interest. We need a common understanding of the new challenges and we need to design new models, new algorithms, new database to count for the user-centric requirements of emerging applications.

The **PersDB 2010 workshop** aims at providing a forum for presentation of the latest research results, new technology developments, and new applications in the areas of personalized/socialized access, profile management, and context awareness in database systems.

7th International Workshop on Data Management for Sensor Networks (DMSN'10)

http://www.cs.ucy.ac.cy/~dmsn10/
September 13, 2010
Location:

The workshop's scope includes all important aspects of sensor data management, including data acquisition, processing, and storage in remote wireless networks; the handling of uncertain sensor data; and the management of heterogeneous and sometimes sensitive sensor data in databases. The resource-constrained, lossy, noisy, distributed, and remote nature of wireless sensor networks implies that traditional database techniques often cannot be applied without significant retooling. Challenges associated with acquiring, processing, and archiving large-scale, heterogeneous sets of live sensor data also call for novel data management techniques. The inherently incomplete and noisy nature of sensor data further calls for techniques for data cleaning, inference, approximation. Finally, in many applications, the collecting of sensor data raises important privacy and security concerns that require new protection and anonymization techniques.

4th International Workshop on Management of Uncertain Data (MUD'10)

http://ewi1276.ewi.utwente.nl:3000/
September 13, 2010
Location:

The aim of the workshop on Management of Uncertain Data is to provide a forum to share original ideas as well as research results and practical development experiences among researchers and application developers. In this workshop we want to explore the various aspects of uncertainty in data as well as techniques how to handle them in the domain of databases. In particular, we are interested in discussing the different kinds of uncertainty, different models for uncertainty representation in databases, techniques for querying and updating data involving uncertainty, and the varying application areas in which handling uncertain data is involved.

1st International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS'10)

http://www.adms-conf.org/
September 13, 2010
Location:

The objective of this one-day workshop is to investigate opportunities in accelerating data management systems and workloads (which include traditional OLTP, data warehousing/OLAP, ETL, Streaming/Realtime, and XML/RDF Processing) using CPU (e.g., commodity and specialized Multi-core, Many-core, GPUs, and FPGAs), storage systems (e.g., Storage-class Memories like SSDs and Phase-change Memory), and multicore programming strategies like OpenCL.

The current data management scenario is characterized by the following trends: traditional OLTP and OLAP/data warehousing systems are being used for increasing complex workloads (e.g., Petabyte of data, complex queries under real-time constraints, etc.); applications are becoming far more distributed, often consisting of different data processing components; non-traditional domains such as bio-informatics, social

networking, mobile computing, sensor applications, gaming are generating growing quantities of data of different types; economical and energy constraints are leading to greater consolidation and virtualization of resources; and analyzing vast quantities of complex data is becoming more important than traditional transactional processing.

At the same time, there have been tremendous improvements in the CPU and memory technologies. Newer processors are more capable in the CPU and memory capabilities and are optimized for multiple application domains. Commodity systems are increasingly using multi-core processors with more than 4 cores per chip and enterprise-class systems are using processors with at least 32 cores per chip. Specialized multi-core processors such as the Cell and GPUs have brought the computational capabilities of supercomputers to cheaper commodity machines. On the storage front, FLASH-based solid state devices (SSDs) are becoming smaller in size, cheaper in price, and larger in capacity. Exotic technologies like Phase-change memory are on the near-term horizon and can be game-changers in the way data is stored and processed.

In spite of the trends, currently there is limited usage of these technologies in data management domain. Naive usage of multi-core processors or SSDs often leads to unbalanced system. It is therefore important to evaluate applications in a holistic manner to ensure effective utilization of CPU and memory resources. This workshop aims to understand impact of modern hardware technologies on accelerating core components of data management workloads. Specifically, the workshop hopes to explore the interplay between overall system design, core algorithms, query optimization strategies, programming approaches, performance modelling and evaluation, etc., from the perspective of data management applications

4th International Workshop on Enabling Real-Time Business Intelligence (BIRTE'10)

http://bird.cs.tu-berlin.de:8008/birte2010/
September 13, 2010
Location:

In today's competitive and highly dynamic environment, analyzing data to understand how the business is performing and to predict outcomes and trends has become critial. The traditional approach to reporting is no longer adequate, users now demand easy-to-use intelligent platforms and applications capable of analyzing real-time data to provide insight and actionable information at the right time. The end goal is to support better and timelier decision making, enabled by the availability of up-to-date, high quality information.

Although there has been progress in this direction and many companies are introducing products towards meeting this goal, there is still a long way to go. In particular, the whole lifecycle of business intelligence requires new techniques and methodologies capable of dealing with the new requirements imposed by the new generation of BI applications. From the capturing of real-time business data to the transformation and delivery of actionable information, all the stages of the Business Intelligence (BI) cycle call for new algorithms and paradigms as the basis of new functionalities including dynamic integration of real-time data feeds from operational sources, optimization and evolution of ETL transformations and analytical models, and dynamic generation of adaptive real-time dashboards, just to name a few.

Workshop on Semantic Data Management (SemData'10)

http://semdata.org/events/2010/vldb
September 17, 2010
Location:

The Semantic Web represents the next generation Web of Data, where information is published and interlinked in order to facilitate the exploitation of its structure and meaning for both humans and machines. Semantic Web applications require database management systems for the handling of structured data, taking into consideration the models used to represent semantics. To foster the realization of the Semantic Web, the World Wide Web Consortium (W3C) developed a set of metadata models, ontology models, and query languages. Today, most of the Semantic Web repositories are database engines, which store data represented in RDF, support SPARQL queries, and can interpret schemata and ontologies represented in RDFS and OWL. We are thus at the point where the adoption of semantic technologies is growing. However, these technologies often appear to be immature, and tend to be too expensive or risky to deploy in real business. Solid data

management layer concepts, architectures, and tools are important to everyone in the semantic ecosystem, and creating them requires a strong community, with a critical mass of involvement.

Semantic data management refers to a range of techniques for the manipulation and usage of data based on its meaning. It enables sustainable solutions for a range of IT environments, where the usage of today's mainstream technology is either inefficient or entirely unfeasible: enterprise data integration, life science research, data sharing in SaaS architectures, querying linked data on the Web. In a nutshell, semantic data management fosters the economy of knowledge, facilitating more comprehensive usage of larger scale and more complex datasets at lower cost.

The goal of the SemData workshop is to provide a platform for the discussion and investigation of various aspects related to semantic databases and data management in the large. Many of the semantic data management challenges cumulate in the need for scalable and performing database solutions for semantic data, a building block that runs largely behind comparable non-semantic technologies. In order to make semantic technologies take on the targeted market share, it is indispensable that technological progress allows semantic repositories to reach near performance parity with some of the best RDBMS solutions without having to omit the advantages of a higher query expressivity compared to basic key-value stores, or the higher schema flexibility compared to the relational model. It is time that one must no longer pay a heavy price in terms of longer run times or more expensive equipment for profiting from the flexibility of the generic physical model underlying the semantic graph-based structures of RDF. We also recognize that there will always be a burden with more flexibility. Hence, the goal is to minimize the drawbacks and maximize the advantages of the semantic RDF-minded repositories.

7th VLDB Workshop on Secure Data Management (SDM'10)

http://www.hitech-projects.com/sdm-workshop/sdm10.html

September 17, 2010

Location:

Although cryptography and security techniques have been around for quite some time, emerging technologies such as ubiquitous computing and ambient intelligence that exploit increasingly interconnected networks, mobility and personalization, put new requirements on security with respect to data management. As data is accessible anytime anywhere, according to these new concepts, it becomes much easier to get unauthorized data access. Furthermore, it becomes simpler to collect, store, and search personal information and endanger people's privacy. Therefore, research in the area of secure data management is of growing importance, attracting attention of both the data management and security research communities The interesting problems range from traditional ones such as, access control (with all variations, like dynamic, context-aware, role-based), database security (e.g. efficient database encryption schemes, search over encrypted data, etc.), privacy preserving data mining to controlled sharing of data.

The aim of the workshop is to bring together people from the security research community and data management research community in order to exchange ideas on the secure management of data. This year an additional special session will be organized with the focus on secure and private data management in healthcare. Data security and privacy issue are traditionally important in the medical domain. However, recent developments and increasing deployment of IT in healthcare such as the introduction of electronic health records and extramural applications in the personal health care domain, pose new challenges towards the protection of medical data. In contrast to other domains, such as financial, which can absorb the cost of the abuse of the system, healthcare cannot. Once sensitive information about individual's health problems is uncovered and social damage is done, there is no way to revoke the information or to restitute the individual. In addition to this, the medical field has some other specific characteristics, such as long-term value of medical data and flexibility with respect to, on one hand confidentiality, and on the other hand availability of medical data in the case of emergency.

7th International XML Database Symposium (XSym'10)

http://www.lirmm.fr/~bella/XSym/
September 17, 2010
Location:

The XML Database Symposium (XSym) series focuses on the convergence of database technology with XML technology, and brings together academics, practitioners, users and vendors to discuss the use and synergy between these technologies.

2nd TPC Technology Conference on Performance Evaluation & Benchmarking (TPCTC'10)

http://www.lirmm.fr/~bella/XSym/
September 17, 2010
Location:

First established in August 1988, the TPC has shaped the landscape of modern transaction processing and database benchmarks over two decades. Now, the world is in the midst of an extraordinary information explosion. Enterprise data and user generated data levels continue to grow exponentially. This has challenged researchers and industry experts to develop innovative techniques to evaluate and benchmark software and hardware technologies. As a result, the TPC is conducting its second Technology Conference (TPCTC 2010) in conjunction with the 36th International Conference on Very Large Data Bases (VLDB 2010).

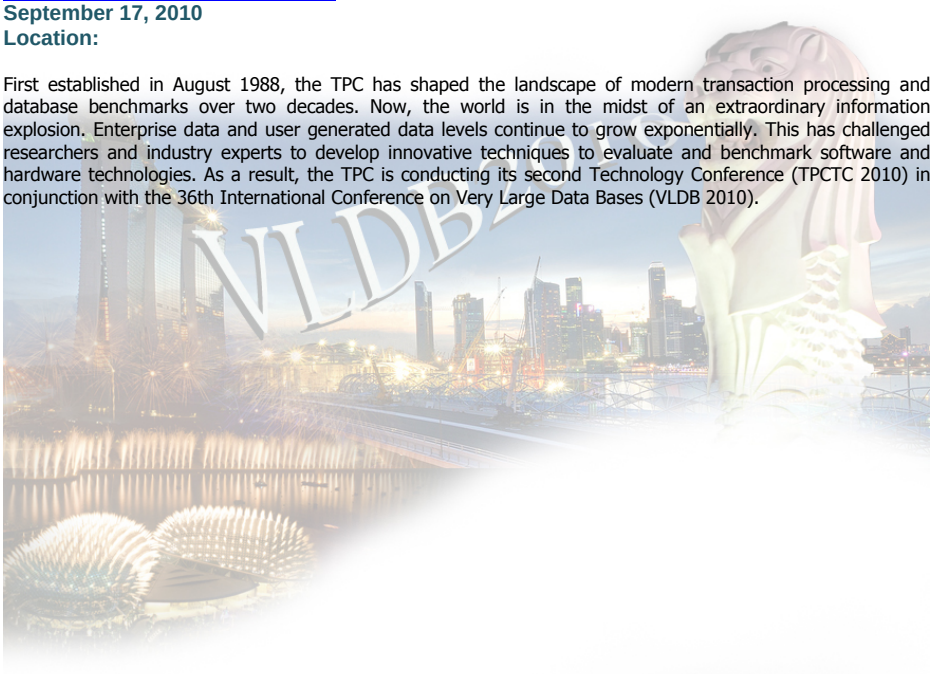

VLDB ORGANIZATION

General Co-Chairs
Tok Wang LING, National University of Singapore, Singapore
Krithi RAMAMRITHAM, Indian Institute of Technology Bombay, India

Technical Program Chair
Elisa BERTINO, Purdue University, USA

Core Database Technology Program Chair
Kian Lee TAN, National University of Singapore, Singapore

Infrastructure for Information Systems Program Chair
Paolo ATZENI, Università Roma Tre, Italy

Industrial Applications and Experience Program Chair
Philip S. YU, University of Illinois at Chicago, USA

Experiments and Analyses Program Chair
Masaru KITSUREGAWA, University of Tokyo, Japan

Demonstrations Program Co-Chairs
Ralf Hartmut GÜTING, University of Hagen, Germany
Wang-Chien LEE, The Pennsylvania State University, USA
Ee-Peng LIM, Singapore Management University, Singapore
Xiaofeng MENG, Remin University, China

PhD Program Co-Chairs
John MYLOPOULOS, University of Toronto, Canada
Lizhu ZHOU, Tsinghua University, China
Xiaofang ZHOU, University of Queensland, Australia

Workshop Program Co-Chairs
Amol DESHPANDE, University of Maryland, USA
Zachary G. IVES, University of Pennsylvania, USA
Anthony Kum Hoe TUNG, National University of Singapore, Singapore

Tutorial Program Co-Chairs
Guozhu DONG, Wright State University, USA
Hwee Hwa PANG, Singapore Management University, Singapore

Panel Program Co-Chairs
Marianne WINSLETT, University of Illinois at Urbana-Champaign, USA
Jeffrey Xu YU, The Chinese University of Hong Kong, China

Ten-Year Best Paper Award Committee
Stefano CERI, Politecnico di Milano, Italy
M. Tamer ÖZSU, University of Waterloo, Canada (Chair)
Kyu-Young WHANG, KAIST, South Korea

Best Paper Award Committee
Paolo ATZENI, Università Roma Tre, Italy
Elisa BERTINO, Purdue University, USA (Chair)
Christian S. JENSEN, Aalborg University, Denmark
Meral ÖZSOYOĞLU, Case Western Reserve University, USA
Kian-Lee TAN, National University of Singapore, Singapore
Yufei TAO, Chinese University of Hong Kong, China

VLDB Endowment Liaison
David LOMET, Microsoft Research, USA

Geographical Area Coordinators

(Americas)
Jiawei HAN, University of Illinois at Urbana-Champaign, USA

(Asia-Pacific)
Kyu-Young WHANG, KAIST, South Korea

(Europe, Mideast, and Africa)
Timos SELLIS, National Tech. University of Athens, Greece

National Database Society Liaisons, Asia Pacific

(CCF-DBS)
Jianzhong LI, Harbin Institute of Technology, China,

(DB India)
Srinath SRINIVASA, IIIT Bangalore, India

(The Database Society of Japan (DBSJ))
Haruo YOKOTA, Tokyo Institute of Technology, Japan

(KIISE (Korean Institute of Information Scientists and Engineers) Database Society)
Bong Hee HONG, Pusan National University, South Korea

Local Arrangement Co-Chairs

Sourav S BHOWMICK, Nanyang Technological University, Singapore
Mong Li LEE, National University of Singapore, Singapore

Sponsorships Co-Chair

Limsoon WONG, National University of Singapore, Singapore
Dan SUCIU, University of Washington, USA

Publicity Co-Chairs

Stéphane BRESSAN, National University of Singapore, Singapore
Umeshwar DAYAL, HP Labs, USA
Christian S. JENSEN, Aalborg University, Denmark

Proceedings Co-Chairs

Yi CHEN, Arizona State University, USA
Y.C. TAY, National University of Singapore, Singapore

Registration Chair

Chee Yong CHAN, National University of Singapore, Singapore

Treasurer

Wynne HSU, National University of Singapore, Singapore

Web Management

Stéphane BRESSAN, National University of Singapore, Singapore
Zaini MOHD, National University of Singapore, Singapore

Conference Secretariat

Siew Foong HO, National University of Singapore, Singapore

Core Database Technology

Kian Lee TAN, **Program Chair** (National University of Singapore, Singapore)
Daniel ABADI (Yale University, USA)
Ashraf ABOULNAGA (University of Waterloo, Canada)
Laurent AMSALEG (IRISA-CNRS, France)
Walid AREF (Purdue University, USA)
Sourav S BHOWMICK (Nanyang Technological University, Singapore)
Angela BONIFATI (Icar-CNR, Italy)
Peter BONCZ (CWI, The Netherlands)
Philippe BONNET (University of Copenhagen, Denmark)
Luc BOUGANIM (INRIA, France)
Stéphane BRESSAN (National University of Singapore, Singapore)
Nicolas BRUNO (Microsoft Research, USA)
Barbara CATANIA (Universita di Genova, Italy)
Chee Yong CHAN (National University of Singapore, Singapore)
Ugur CETINTEMEL (Brown University, USA)
Lei CHEN (Hong Kong University of Science and Technology, China)
Shimin CHEN (Intel Research Pittsburgh, USA)
Reynold CHENG (University of Hong Kong, China)
Brian COOPER (Yahoo! Research, USA)
Bin CUI (Peking University, China)
Gautam DAS (University of Texas, Arlington, USA)
Amol DESHPANDE (University of Maryland, USA)
Alin DEUTSCH (University of California, San Diego, USA)
Yanlei DIAO (University of Massachusetts, Amherst, USA)
Wenfei FAN (University of Edinburg, UK)
Alan FEKETE (University of Sydney, Australia)
Elena FERRARI (University of Insubria, Italy)
Shel FINKELSTEIN (SAP, USA)
Peter FISCHER (ETH Zurich, Switzerland)
Minos GAROFALAKIS (Technical University of Crete, Greece)
Johannes GEHRKE (Cornell University, USA)
Gabriel GHINITA (Purdue University)
Leo GIAKOUMAKIS (Microsoft, USA)
Torsten GRUST (Technische Universität München, Germany)
Jarek GRYZ (York University, USA)
Wook-Shin HAN (Kyungpook National University, Korea)
Jayant HARITSA (Indian Institute of Science, India)
Stavros HARIZOPOULOS (HP Labs, Palo Alto, CA)
Zachary IVES (University of Pennsylvania, USA)
Christian S. JENSEN (Aalborg University, Denmark)
Chris JEMAINE (University of Florida / Rice University, USA)
Panos KALNIS (KAUST, Saudi Arabia)
Panagiotis KARRAS (National University of Singapore)
Murat KANTARCIOGLU (University of Texas, Dallas, USA)
Alfons KEMPER (Technische Universität München, Germany)
George KOLLIOS (Boston University, USA)
Yannis KOTIDIS (Athens University of Economics and Business, Greece)
Ramamohanarao KOTAGIRI (University of Melbourne, Australia)
Nick KOUDAS (University of Toronto, Canada)
Laks LAKSHMANAN (The University of British Columbia, Canada)
Chen LI (University of California, Irvine & BiMaple)
Jianzhong LI (Harbin Institute of Technology, China)
Lipyeow LIM (IBM T J Watson Research Center, USA)
David LOMET (Microsoft Research, USA)
Boon Thau LOO (University of Pennsylvania, USA)
Qiong LUO (Hong Kong University of Science and Technology, China)
Sharad MEHROTRA (University of California, Irvine, USA)
Sergey MELNIK (Google, USA)
Gerome MIKLAU (University of Massachusetts, Amherst, USA)
Guido MOERKOTTE (University of Mannheim, Germany)

Bongki MOON (University of Arizona, USA)
Mario NASCIMENTO (University of Alberta, Canada)
Wolfgang NEJDL (University of Hannover, Germany)
Christopher OLSTON (Yahoo! Research, USA)
M. Tamer ÖZSU (University of Waterloo, Canada)
HweeHwa PANG (Singapore Management University, Singapore)
Dimitris PAPADIAS (Hong Kong University of Science and Technology, China)
Evaggelia PITOURA (University of Ioannina, Greece)
Neoklis POLYZOTIS (University of California, Santa Cruz, USA)
Rachel POTTINGER (The University of British Columbia, Canada)
Sunil PRABHAKAR (Purdue University, USA)
Rajeev RASTOGI (Yahoo! Research, India)
Tore RISCH (Uppsala University, Sweden)
Kenneth ROSS (Columbia University, USA)
Elke RUNDENSTEINER (Worcester Polytechnic Institute, USA)
Kai-Uwe SATTLER (Technische Universität Ilmenau, Germany)
Ralf SCHENKEL (Max Planck Institute, Germany)
Eugene SHEKITA (IBM Almaden Research Center, USA)
Hengtao SHEN (University of Queensland, Australia)
Kyuseok SHIM (Seoul National University, Korea)
Richard SNODGRASS (University of Arizona, USA)
Divesh SRIVASTAVA (AT&T, USA)
Ioana STANOI (IBM T.J. Watson Research Center, USA)
Dan SUCIU (University of Washington, USA)
S. SUDARSHAN (IIT Bombay, India)
Wang-Chiew TAN (University of California, Santa Cruz, USA)
Yufei TAO (Chinese University of Hong Kong, China)
Victor VIANU (University of California, San Diego, USA)
Gerhard WEIKUM (Max-Planck Institute, Germany)
Jun YANG (Duke University, USA)
Carlo ZANIOLO (University of California, Los Angeles, USA)
Donghui ZHANG (Microsoft, USA)
Sai WU (National University of Singapore, Singapore)
Aoying ZHOU (East China Normal University, China)
Yongluan ZHOU (University of Southern Denmark, Denmark)

Infrastructure for Information Systems

Paolo ATZENI, **Program Chair** (Università Roma Tre, Italy)
Karl ABERER (EPFL, Switzerland)
Foto AFRATI (NTUA, Greece)
Rakesh AGRAWAL (Microsoft, USA)
Suad ALAGIC (University of Southern Maine, USA)
Mike ATALLAH (Purdue University, USA)
Ricardo BAEZA-YATES (Yahoo! Research, Spain)
Elena BARALIS (Politecnico di Torino, Italy)
Denilson BARBOSA (University of Alberta, Canada)
Leopoldo BERTOSSI (Carleton University, Canada)
Boualem BENATALLAH (University of New South Wales, Australia)
Jan Van den BUSSCHE (Universiteit Hasselt, Belgium)
Jose BLAKELEY (Microsoft, USA)
Mike CAFARELLA (University of Washington, USA)
Michael CAREY (University of California, Irvine, USA)
Sharma CHAKRAVARTHY (The University of Texas at Arlington, USA)
Kevin CHANG (University of Illinois at Urbana-Champaign, USA)
Arbee CHEN (National Chengchi University, Taiwan)
Laura CHITICARIU (IBM Research, USA)
Jan CHOMICKI (University at Buffalo, USA)
Chris CLIFTON (Purdue University, USA)
Susan DAVIDSON (University of Pennsylvania, USA)
Umeshwar DAYAL (HP Labs, USA)
Lois DELCAMBRE (Portland State University, USA)
AnHai DOAN (University of Wisconsin, USA)
Jens DITTRICH (Saarland University, Germany)
Michael FRANKLIN (University of California, Berkeley, USA)
Johann-Christoph FREYTAG (Humboldt-Universität zu Berlin, Germany)
Aristides GIONIS (Yahoo! Research, Spain)
Jonathan GOLDSTEIN (Microsoft, USA)
Philip GOLLE (PARC, USA)
Luis GRAVANO (Columbia University, USA)
Ralf Hartmut GÜTING (Fernuniversität Hagen)
Alon HALEVY (Google, Inc., USA)
Mauricio HERNANDEZ (IBM Research, USA)
Howard HO (IBM Research, USA)
Rick HULL (IBM Research)
Panos IPEIROTIS (New York University, USA)
Anastasios KEMENTSIETSIDIS (IBM Research, USA)
Maurice van KEULEN (University of Twente, The Netherlands)
Christoph KOCH (Cornell University, USA)
Hank KORTH (Lehigh University, USA)
Donald KOSSMANN (ETH Zürich, Switzerland)
Georgia KOUTRIKA (Stanford University, USA)
Peer KROEGER (Ludwig-Maximilians-Universitaet München, Germany)
Alexandros LABRINIDIS (University of Pittsburgh, USA)
Alberto LAENDER (Federal University of Minas Gerais, Brazil)
Maurizio LENZERINI (Sapienza, Università di Roma, Italy)
Frederick LOCHOVSKY (HKUST, Hong Kong)
Jiaheng LU (Renmin University of China)
Ioana MANOLESCU (INRIA, France)
Yannis MANOLOPOULOS (Aristotle University, Thessaloniki, Greece)
Giansalvatore MECCA (Università della Basilicata, Italy)
Paolo MERIALDO (Università Roma Tre, Italy)
Moira NORRIE (ETH Zuerich, Switzerland)
Z. Meral ÖZSOYOĞLU (Case Western Reserve University, USA)
Themis PALPANAS (Università di Trento, Italy)
Jan PAREDAENS (University of Antwerp, Belgium)
Norman PATON (University of Manchester, UK)
Hamid PIRAHESH (IBM Research, USA)

Lucian POPA (IBM Research, USA)
Alexandra POULOVASSILIS (Birkbeck, University of London, UK)
Erhard RAHM (Universität Leipzig, Germany)
Matthias RENZ (Ludwig-Maximilians-Universität München, Germany)
Marc SCHOLL (Universität Konstanz)
Ning SHANG (Purdue University, USA)
Jayavel SHANMUGASUNDARAM (Yahoo! Research, USA)
Dennis SHASHA (New York University, USA)
Oded SHMUELI (Technion, Israel)
Jianwen SU (University of California, Santa Barbara, USA)
Val TANNEN (University of Pennsylvania, USA)
Jens TEUBNER (ETH, Switzerland)
Yannis VELEGRAKIS (Università di Trento, Italy)
Gottfried VOSSEN (WWU Muenster, Germany)
Ouri WOLFSON (University of Illinois at Chicago)
Peter WOOD (Birkbeck, University of London, UK)
Cong YU (Yahoo! Research, USA)
Stan ZDONIK (Brown University, USA)
Pavel ZEZULA (Masaryk University, Czech Repubilc)
Zhenjie ZHANG (National University of Singapore, Singapore)

Industrial Applications and Experience

Philip S. YU, **Program Chair** (University of Illinois at Chicago, USA)
Henrique ANDRADE (IBM USA)
Michael BRODIE (Verizon, USA)
Malu CASTELLANOS (HP Labs, USA)
Edward CHANG (Google, China)
Yuan-chi CHANG (IBM, USA)
Ming-Syan CHEN (NTU, Taiwan)
Yun CHI (NEC USA)
Graham CORMODE (AT&T Labs, USA)
Takeshi FUKUDA (IBM, Japan)
Bugra GEDIK (IBM, USA)
Alex GORELIK (Exeros, USA)
Goetz GRAEFE (HP Labs, USA)
Hui I HSIAO (IBM USA)
Meichun HSU (HP Labs, China)
Dean JACOBS (SAP, Germany)
Chiao LIU (Microsoft, USA)
Bo LONG (Yahoo!, USA)
Wei-Ying MA (Microsoft, China)
Ravi MURTHY (Oracle, USA)
Makoto ONIZUKA (NTT Labs, Japan)
Alex SZALAY (JHU, USA)
Ming-Chien SHAN (SAP, China)
Evimaria TERZI (IBM USA)
Michale VLACHOS (IBM, Swizterland)
Haixun WANG (Microsoft, China)
Xiaoxin YIN (Microsoft USA)

Experiments and Analyses

Masaru KITSUREGAWA, **Program Chair** (University of Tokyo, Japan)
Kasim CANDAN (ASU, USA)
Umeshwar DAYAL (HP Labs, USA)
H JAGADISH (MU, USA)
Hiroyuki KITAGAWA (Tsukuba University, Japan)
Paul LARSON (Microsoft Research, USA)
Wen-Syan LI (SAP, China)
Xuemin LIN (U of New South Wales .Australia)

Ling LIU (George Tech., USA)
Sanjay MADRIA (UMR, USA)
Volker MARKL (Tu-Nerlin, Germany)
Beng Chin OOI (NUS, Singapore)
Cyrus SHAHABI (USC, USA)
Kyu-Young WHANG (KAIST Korea)
Kung Lung WU (IBM research Watson, USA)
Jefferey YU (CUHK, HongKong)

Demonstrations

Ralf Hartmut GÜTING, **Program Co-Chair** (University of Hagen, Germany)
Wang-Chien LEE, **Program Co-Chair** (The Pennsylvania State University, USA)
Ee-Peng LIM, **Program Co-Chair** (Singapore Management University, Singapore)
Xiaofeng MENG, **Program Co-Chair** (Remin University, China)
Jidong CHEN (EMC China Lab, China)
Lei CHEN (Hong Kong University of Science and Technology, China)
Xin DONG (AT&T Labs Research, USA)
Bugra GEDIK (IBM Watson Research Center, USA)
Fosca GIANNOTTI (Istituto di Scienza e Tecnologie dell'Informazione Alessandro Faedo", Italy)
Hyoil HAN (LeMoyne-Owen College, USA)
Wook-Shin HAN (Kyung-Pook National University, Korea)
San-Yih HWANG (National Sun Yat Sen University, Taiwan)
Daxing JIAN (Microsoft Research Asia, China)
Vana KALOGERAKI (Univ. of California, Riverside, USA)
Daniel KIFER (Pennsylvania State University, USA)
Ulf LESER (Humboldt University, Berlin, Germany)
Weiyi MENG (Binghamton University, USA)
Mohamed F. MOKBEL (University of Minnesota, USA)
Felix NAUMANN (Hasso Plattner Institut, Germany)
Frank NEVEN (Hasselt University, Belgium)
Daniela NICKLAS (Uni Oldenburg, Germany)
Alex NTOULAS (Microsoft Search Labs, USA)
Yasushi SAKURAI (NTT Communication Science Laboratories, Japan)
Peter SCHEUERMANN (Northwestern University, USA)
Katsumi TANAKA (Kyoto University, Japan)
Nesime TATBUL (ETH Zuerich, Switzerland)
Jens TEUBNER (ETH Zürich Switzerland)
Jianyong WANG (Tsinghua University, China)
Li XIONG (Emory University, USA)
Denian YANG (Academic Sinica, Taiwan)
Ge YU (Northeastern University, China)
Baihua ZHENG (Singapore Management University, Singapore)

PhD Workshop

John MYLOPOULOS, **Program Co-Chair** (University of Toronto, Canada)
Lizhu ZHOU, **Program Co-Chair** (Tsinghua University, China)
Xiaofang ZHOU, **Program Co-Chair** (University of Queensland, Australia)
Lei CHEN (Hong Kong University of Science and Technology, China)
Ke DENG (University of Queensland, Australia)
Alan FEKETE (University of Sydney, Australia)
Ihab ILYAS (University of Waterloo, Canada)
Peer KRÖGER(Ludwig-Maximilians-Universität München, Germany)
Xuemin LIN (University of New South Wales, Australia)
Amélie MARIAN (RutgersUniversity, USA)
Miyuki NAKANO (University of Tokyo, Japan)
Yoshiharu ISHIKAWA (Nagoya University, Japan)

Jian PEI (Simon Fraser University, Canada)
Jagan SANKARANARAYANAN (University of Maryland, USA)
Divesh SRIVASTAVA (AT&T, USA)
Fabio SCHREIBER (Politecnico di Milano, Italy)
Mohamed SHARAF (University of Queensland, Australia)
Kyuseok SHIM (National Seoul University, Korea)
Anthony TUNG (National University of Singapore, Singapore)
Yannis VELEGRAKIS (University of Trento, Italy)
Kyu-Young WHANG (KAIST, Korea)
Ke YI (Hong Kong University of Science and Technology, China)
Jeffrey X. YU (Chinese University of Hong Kong, China)
Aoying ZHOU (East China Normal University, China)

SINGAPORE – GATEWAY TO ASIA

Singapore, founded as a British trading colony in 1819, declared its independence in 1965. Over the next four decades, Singapore worked hard to overcome numerous hurdles, attaining an average growth rate of 7.5% per annum.

Located right in the heart of Asia, Singapore thrives as a financial and trading hub for the region and beyond. Through global trade linkages and establishments of Free Trade Areas, ties with China, India and the rest of Asia are continually strengthened. Language fluency, coupled with excellent technological and management expertise gives Singapore an edge in accessing this growing region.

Even as Singapore has transformed into a cosmopolitan city state, it has not lost its Asian charms, reflected in its strong multi-cultural society. This is evident in its hospitality, cuisine, arts, and architecture as Singapore continues to evolve, successfully blending the East and West.

GEOGRAPHY & WEATHER

Singapore is an island country in Southeast Asia. It is about 140 kilometres north of the equator and has an area of 639 square kilometres. Singapore is 42 kilometres long and 23 kilometres across at its widest point. About half of the offshore islands of Singapore are uninhabited. Some of them have been developed for recreation and tourism. Singapore's accessibility by air and sea help consolidate its status as a favourite global convention centre. Singapore's location within South East Asia makes it a favourite tourism destination, as well as a convenient hub for "short-hops" to neighbouring countries.

Singapore has an equatorial climate, which is warm and humid all year round. The temperature averages around 28 degrees Celsius daily, with abundant rainfall during the monsoon seasons from December to March and June to September. In between, the weather is characterized by light and variable surface winds and rain ranging from light showers to thunderstorms.

POPULATION

Singapore has a population of about 4.2 million, of which one-fifth comprise of non-resident workers and expatriates. The resident population is multiethnic comprising 77% Chinese, 14% Malays, 7.6% Indians and 1.4% are of diverse ethnicity. Singapore enjoys racial harmony and social cohesion, and is one of the most cosmopolitan societies in the world, being home to a sizeable expatriate community.

LANGUAGE

English is the language of education and is widely spoken throughout Singapore. It unites people of different races and different tongues, and is used in all shops and places of business. That said, it is by no means uncommon to hear conversation being conducted in Mandarin, Malay, Tamil or one of the many dialects. Ask for help or directions in English, though, and they will switch languages in mid-sentence to help you! Being an international tourist destination, Singapore has a fully-developed tourism industry, and this includes professional guides and translators fluent in a variety of foreign languages. An increasing number of Singaporeans are fluent in French, German and Japanese.

INFRASTRUCTURE AND COMMUNICATIONS

Singapore has a well-developed infrastructure to meet all business needs. The city was ranked 3rd among 61 countries for infrastructure in the World Competitiveness Yearbook 2006. Singapore is also an important telecommunications centre in this region and was ranked 2nd most technologically ready nation in Asia in the Global Competitiveness Report 2006-2007. The first nation-wide broadband information infrastructure in the world is available island wide and Singapore has also established high-speed connectivity to more than 30 countries.

TRAVELLING TO SINGAPORE

Singapore's Changi International Airport is a premier international air hub. A liberal aviation policy, which entails an exchange of traffic rights with other countries, enables airlines to expand connectivity between Singapore and the rest of the world. This has led to Singapore being a crucial point on air traffic routes across the world. There are more than 80 international airlines operating scheduled services through Singapore to 185 cities in 58 countries worldwide. The national carrier, Singapore Airlines, together with its code-share partners, flies to 90 cities in 40 countries. Every week, more than 4,000 scheduled flights provide quick and convenient connections from Singapore to major destinations worldwide.

The Budget Terminal, customized for low cost carriers is able to handle about 2.7 million passengers per annum. As at 1 March 2007, low cost carriers in Singapore operate 566 weekly flights to 27 cities. Terminal 3 opened on 9 January 2008 and is the largest of Changi Airport's three terminals, handling 22 million passengers per annum.

TRAVELING WITHIN SINGAPORE

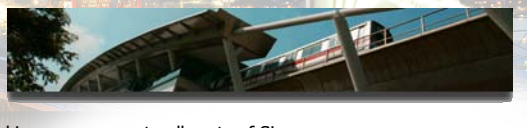

Singapore has a world class, highly efficient and reliable public transport system. It consists of the mass rapid transit system, bus services and taxis. Fast, efficient and comfortable, the Mass Rapid Transit and Light Rapid Transit systems are the backbone of the city's public transport system, taking passengers to all parts of Singapore.

SAFETY

Singapore is one of the safest cities in the world, where both locals and visitors can walk along the streets late at night without fear. The nation's general economic prosperity and professional, effective police force contribute to the low crime rate, and enable continued success in controlling crime and enhancing personal security.

Over the last 15 years, Singapore's security and emergency response agencies have built up their capabilities to respond to contingencies or threats. These plans are in place and are exercised regularly over the years.

Singapore is globally ranked third in the "protection of personal security and private property" and first in "social cohesion as a priority for the government" by the World Competitiveness Yearbook (WCY), published by the International Institute for Management Development. The WCY also identified Singapore as the most politically stable country in Asia. In research done by Mercer Human Resource Consulting in 2005, Singapore was ranked second in personal safety among the world's top 50 cities.

AFTER CONFERENCE HOURS

Uniquely Singapore – "unique" is the word that best captures Singapore - a dynamic, cosmopolitan city-state where different cultures, ethnic groups and religions blend harmoniously. A bridge between the East and the West for centuries, Singapore, located in the heart of fascinating Southeast Asia, continues to embrace tradition and modernity today. Brimming with unbridled energy and bursting with exciting events, the city-state offers countless unique memorable experiences waiting to be discovered.

With such an exciting diversity of ethnic cultures and attractions, there is always something to do in Singapore. The Lion City also offers duty-free shopping, more than 30 tempting local, Asian and international cuisines plus sports galore including a large number of lush golf courses!

Singapore offers exceptional value for money in terms of shopping, dining; entertainment and accommodation. Delegates will enjoy a high quality experience at an affordable price. Furthermore the universal appeal of Singapore, both within Asia and long-haul markets, will ensure maximum delegate attendance at the Conference.

A spectacular array of arts events is offered all year round. Singapore has it all, from the fine art and antique exhibitions, award-winning music and dance extravagances, to cutting-edge drama, film and comedy. Delegates can immerse themselves in the sights and sounds of traditional Chinese opera, embark on a leisurely exploration of Singapore's arts district, visit museums and check out the various performing arts venues and institutions, including the world class arts complex, The Esplanade – Theatres on the Bay, roll the dice at the brand new casinos.

ATTRACTIONS IN SINGAPORE

We are pleased to showcase what Singapore has to offer to VLDB delegates in 2010.

Singapore is a vibrant, multi-cultural, cosmopolitan city that embodies the essence of a world class city. A city rich in contrast and colours, Singapore is a harmonious blend of culture, cuisine, arts and architecture where modernity and tradition juxtapose to create a unique experience. Brimming with unbridled energy and bursting with exciting events, the city offers countless memories waiting to be discovered. There is something for everyone in Singapore, from culture, arts, nature and adventure Singapore style to absolute indulgence. Here are some suggested itineraries available for selection.

CULTURAL HERITAGE

CHINATOWN

No visit to Singapore would be complete without a visit to Chinatown – the original site where Chinese immigrants lived and worked. Chinatown is where delegates can visit a wet market and see housewives haggling for the best fresh produce available or experience a Chinese tea appreciation session with an opera performance. Shop for Chinese antiquities or buy Tiger Balm (medicated ointment) and marvel at the array of Chinese herbs on sale. Visit the famous Chinese temples there or the Sri Mariamman Temple, right in the heart of Chinatown.

The bustling street market, an integral part of Chinatown's street scene in the past, has made a comeback from June 2004. The new Chinatown Night Market has over 200 street stalls selling a wide and interesting selection of retail items, ranging from handicraft to fashion apparel, electronic goods, jewellery, trinkets, modern home accessories and food items, just to name a few.

This district also houses the Chinatown Heritage Centre in three newly restored shophouses. It showcases a wealth of memories and untold stories of the past. Each level of the Centre takes the visitor to a different period in the history of Chinatown and the lives of Singapore's early settlers.

LITTLE INDIA

Your senses will come alive in this "colourful" ethnic area, which used to be just marshland and grazing fields for cattle. Little India, where the first Indian settlers in Singapore resided when they arrived with Sir Stamford Raffles as assistants and soldiers back in the 1800s.

Walk along the five-foot ways of shophouses and discover an array of Indian silk, brassware and carvings. For those who are more adventurous, the mildly-narcotic betel nut, believed to be good for the digestive system, is a must-try. Or have a South Indian rice meal with your fingers off a banana leaf at Apollo Banana Leaf Restaurant or Muthu's Curry Restaurant along Race Course Road.

ARAB STREET / KAMPONG GLAM

Embroidered materials, semi-precious stones, basketware, accessories for a pilgrimage, natural oils and perfumes abound in Arab Street / Kampong Glam**,** which got its name from the Glam trees that once grew there. Medicinal oil was extracted from the tree and its bark used by the Buginese and Malays to caulk their boats. A visit to the Sultan Mosque, the most prominent mosque in the city, is a must.

PERANAKAN HALL

Peranakans refer to the descendents of the early Chinese community who settled in the Malay Archipelago from the 17th century. Its culture is a rich blend of the Chinese and Malay cultures, with some influence from the Portuguese, Dutch, British, Thai, Indian and Indonesian cultures.

The tour will start off at Katong, the Peranakan enclave of Singapore. Katong is well known as a food haven among the locals. As a Peranakan enclave, it is a district that is rich in Peranakan and Eurasian heritage and architecture. Many of the older Peranakan homes, now standing back from the sea, were once seaside villas.

Wander along East Coast Road for a choice of delectable local fare, including the famous laksa dish, tasty local seafood and Peranakan cuisine.

RELIGIOUS HERITAGE

First stop at the Thian Hock Keng Temple (one of the oldest Chinese Temples in Singapore), built in 1821 by seamen from China who were grateful to the gods for their safe passage. The unique wooden structures and wood carvings were built and assembled without using any nails.

Another architectural treasure is the Sri Mariamman Temple, Singapore's oldest Hindu temple that is situated in Chinatown. Dating back to 1827, the brick bones of the present building were built around 1843. Later additions have included a profusion of deities carved all over its walls, bell-decked doors and frescoes on the ceilings.

The last stop is the Sultan Mosque. Named after the first Sultan of Singapore, the imposing structure was built with a SGD3,000 grant from the East India Company. This is the largest mosque in Singapore, and the present building was constructed in 1928. With its massive gold-topped dome and spacious prayer hall, the mosque has one particularly interesting feature – the base of the dome is made up of glass bottles.

OLD WORLD CHARM

Boat Quay Route

Retrace Singapore's transformation from a fishing village into a bustling metropolis. Visit Sir Stamford Raffles landing site, the Parliament House, Victoria Theatre and Empress Place, all of which are 19th Century sites that have been restored.

Enjoy a leisurely bumboat ride down the Singapore River, which used to be the lifeline of Singapore. It was here that early immigrants eked out a meager living and saw Singapore transform from an obscure little fishing village to a world-renowned seaport. The bumboat ride will be like a travel through time in the island's history.

ARTS CITY

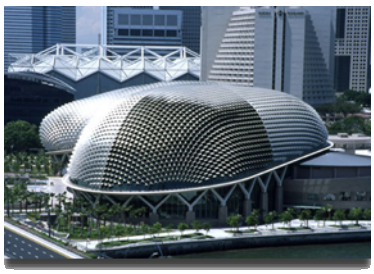

Singapore is also a City for the Arts. It is a burgeoning arts scene and there are drama, music and dance performances all year round at various locations, culminating in the Singapore Arts Festival in June. World-class acts like the WOMAD (World of Music, Art and Dance) have kept their audiences on their toes while the Comedy Festival has sent everyone home in stitches.

FOR THE LOVE OF ART

Calling all art lovers, the Singapore Art Museum (SAM) is a must-see. Transformed from a 19th century mission school building into a state-of-the-art facility for the visual arts, the architecture of the Art Museum melds past and present in a modern space worthy of art exhibitions and programs for today's audience.

The Museum houses the national art collection of Singapore, the largest public collection of 20th-Century Southeast Asian art internationally. Catch highlights from the Museum's permanent collection in Art of Our Time and Beneath The Pavement. Enjoy a cuppa at our cafes or browse through our gift shop for unique gift ideas.

CHINESE OPERA – SIGHTS & SOUNDS

Visit the Chinese Opera Teahouse and discover what the painted faces, the magnificently-embroidered costumes and the brilliantly-colored headgear represent.

Join the experts as they explain the intricacies of Chinese Opera and make-up, and the meaning behind the swish of a sleeve, the flick of a finger or the rhythmic movement of the torso. Showtime starts with a talk-cum-demonstration on Chinese Opera (in English), followed by a colorful performance of an opera excerpt (with English subtitles).

WHERE ASIAN CULTURES COMES ALIVE!

Trace the ancestral heritage of Singapore at the Asian Civilisations Museum, which is housed in the grand Empress Place Building beside the Singapore River. Its ten thematic galleries has over 1300 artefacts from the civilizations of Southeast Asia, South Asia, China and the Islamic world of West Asia. Complementing the displays are state-of-the-art multimedia displays and interactive zones which provide a memorable experience for the senses. The museum has a special exhibition on Japanese masks.

Not to be missed is the Singapore River Interpretive Gallery, which tells the moving story of immigrants who once lived and worked along the Singapore River.

NATURE

TROPICAL SPLENDOUR

Established in 1859, Singapore Botanic Gardens was where the first rubber trees were cultivated from seedlings sent from Kew Gardens in London. Spread over 52 hectares, the garden is a combination of untouched primary forest and specialty gardens displaying palms, orchids, cycads, ginger, bonsai and heritage trees.

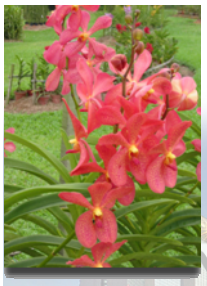

Be charmed at the National Orchid Garden, which has over 1,000 orchid species and 2,000 orchid hybrids. It houses sprawling orchid displays, water features and an exotic bromeliad collection from Central and South America. Expect a pleasant morning of strolling through the gardens with specialist nature guides, who will explain the diversity and history of the garden and the many facts of orchid-growing, as well as special orchid hybrids named after VIPs.

Don't miss the new Cool House – the first conservatory in Asia to simulate a tropical mountain forest. The use of realistic landscapes, rushing torrents, drifting mist, plants and climate-control offers a naturalistic display of high elevation sites in the tropics. Trees and rocks are draped with orchids and carnivorous plants from the cooler tropical highlands.

JUNGLE EXPOSE

The Singapore Zoological Gardens, an open-concept zoo is home to more than 2,000 creatures. The zoo has attracted international acclaim because of its clever use of rock walls and streams as natural barriers. One of its highlights is the world's largest social colony of orang-utans from Borneo and Sumatra. Also, discover how the Zoo breeds and maintains butterflies, scorpions, stick insects and other invertebrates for display at the breeding complex of the Fragile Forest.

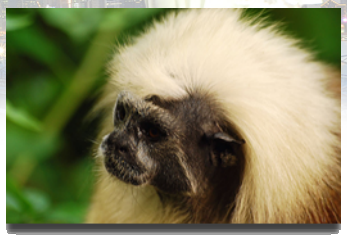

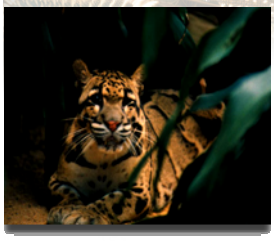

Next to the Zoo is the Night Safari, the world's first night wildlife park set in 40 hectares of dense secondary forest. You will be able to see more than 1,000 nocturnal animals from over 100 species in their natural habitat. As you walk along the well-lit trails, do not be alarmed if a friendly deer walks up to you. The nocturnal activities of animals such as the Cape buffalo, African bongo, blue sheep, striped hyena and the golden jackal will be revealed as you watch them prowling, hunting, feeding and interacting with each other.

Visit the largest open-concept bird park, Jurong Bird Park, in the Asia Pacific and one of the best in the world.

Not-to-be missed attractions in the park include the interesting Penguin Parade, where penguins live and play in an environment similar to that of the South Pole or the South-east Asian Bird Aviary where a thunderstorm is simulated everyday at noon. Take flight only after having refreshments at the "Lodge on Flamingo Lake".

TREASURES IN THE RAINFOREST

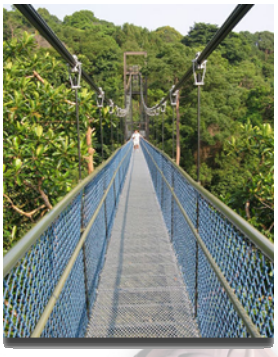

Step into the Central Catchment Nature Reserve and journey back in time to the 19th century, when much of Singapore was covered with lowland, tropical forest. The first of its kind in Singapore, take a walk along the HSBC TreeTop Walk, a 25-metre tall, free-standing suspension bridge that connects the two highest points – Bukit Peirce and Bukit Kalang.

From this vantage point, you can enjoy a panoramic view of the canopy of the rainforest where a rich diversity of flora and fauna dwell resides. See the myriad of interesting animals and insects that on the tree tops, and be totally immersed in the beauty and serenity that surrounds you.

MANGROVES & WETLAND WILDLIFE

The Sungei Buloh Wetland Reserve is an ASEAN (Association of South East Asian Nations) Heritage Park and an internationally-important site for migratory shorebird conservation. The reserve comprises 130 hectares of mangroves, mudflats, ponds and secondary forest.

See the fantastic root formations of mangrove trees, crabs that climb trees, fishes that skip about on land, and metre-long monitor lizards. Spot wild otters hunting in the river and maybe a crocodile or two; all from the safety of boardwalks and walking trails.

LEISURE

SENTOSA ESCAPE

Sentosa is one of the most enchanting Southern islands in Singapore. Originally a fishing village, the island was a British military fortress until 1967, when it was handed back to the Singapore Government.

Today, Sentosa has been transformed into a one-stop multi-million-dollar holiday resort-island with excellent beach amenities and sightseeing opportunities. Visit the Underwater World, Asia's largest tropical oceanarium. Explore the mysteries of ocean life as you enter the depths of this aquatic wonderland and come face-to-face with over 2,300 marine specimens. For those who are truly adventurous, you can even opt to dive with the sharks in the ocenarium. At the Images of Singapore, legend, fact and folklore are creatively interwoven for guests to enjoy a multi-cultural experience. Interact with our walkabout characters and be sure to catch up with them to make your Images of Singapore experience come alive!

RETAIL

Singapore is one place where you can literally "shop till you drop!" The bargains that you get, the variety and quality of merchandise will leave you spoilt for choice. No one who has been to Singapore will forget the Orchard Road shopping experience. The department stores, shopping malls and exclusive boutiques offer a range of products of international quality.

Most malls are interlinked and you can shop in air-conditioned comfort, and not be bothered by tropical showers. If you are adventurous, hop onto an MRT train that takes you to the heartlands, where suburban malls offer equally good quality products at even better prices.

NIGHTLIFE

After hours, Singapore practically transforms itself from an ever-so-efficient business hub to a buzzing network of bars and nightclubs. Busy executives loosen their ties and kick up their heels to the sound of jazz swingin 'at Harry's Bar along Boat Quay, and earnest students-by-day let down their hair for some laid-back bar-hopping at famous Clark Quay.

Pubbing and clubbing at Singapore is a must if you want to feel the night-pulse of the city; experience the city only by day and you get just half of the picture.

Clarke Quay and Riverside Nightlife: Boat Quay and Clarke Quay are dotted with restaurants as well as pubs, so the mood here is slightly laid-back and the crowd mixed. Patronised by serious party-goers, but also people just out for a pint after work, couples chatting over a leisurely dinner, and tourists. The pubs are old Chinese godowns and shophouses given a heartening facelift.

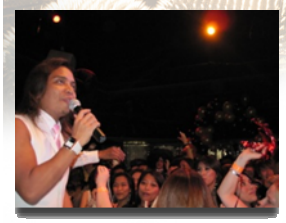

Orchard Road Nightlife: This famous shopping belt also harbours some great clubs. For a touch of soft sophistication, go to Brix, which is situated in the luxurious Grand Hyatt Hotel. The underground place is lit up with candles, and the décor is warm and woody. The pub's resident band plays every day, and the place also has a great wine bar.

Sentosa - Harbour Front Nightlife: For those who like a sea breeze while sipping on a cocktail there's nothing quite like the happy go lucky seaside location of Sentosa and the best bars take full advantage of this fact. So relax and let the sand in between your toes as you forget the stresses of life in this exotic destination.

Civic District Nightlife: Chic bars, elegant bars, ultra-cool clubs and suave restaurants are what are on offer in the Civic District. Stay on after work and meet people like yourself who appreciate ths finer things in life rather than running home at the end of the day.

Marina Bay Nightlife: For cooling, refreshing breezes check out Marina Bay, an oasis of chic that's only really appreciated if you slow down and take your time to explore it. It's a new area with new ideas and it gives you a chance to start anew.

Central Business District (CBD) Nightlife: The Central Business District or CBD is at the heart of this city of over four-and-a-half million people and yes, you need to get away from the crowds occasionally so where better than in a snug bar or top-class restaurant? Consisting roughly of Orchard Road, the Riverside and a chunk of Chinatown the CBD is dotted with such haunts.

Holland Village - Dempsey Hill Nightlife: Holland Village is one of the island-state's most popular nightlife venues and is teeming with restaurants, bars and nightclubs. For this reason Holland Villa is extremely popular with Singapore's younger set and the island's expatriates. It's a fact that you can get almost any sort of cuisine in Holland Village and if you're looking for new bars in which to drink you'll be very busy.

Mohammed Sultan Road Nightlife: In the past few years, this once-forgotten back alley has taken on legendary proportions. Once a ramshackle row of old Chinese godowns, the street has been transformed into a vital, throbbing artery of Singapore's nightlife. The godowns have retained their colourful Peranakan facades, and many of the bars in this area have adopted a Sha nghai-Old World type of interior.

Raffles Hotel: No list of Singapore nightspots would be complete without a mention of the Long Bar at Raffles Hotel. Housed in the historic Raffles Hotel, this is where the Singapore Sling, the island's most famous cocktail, was invented (you can get free serving of this drink if you take Singapore Airlines on your way to VLDB ☺). The décor, in line with that of the hotel, is high-class colonial. Old-style fans flap about on the ceiling and you're encouraged to throw your peanut shells on the floor -- they're there to add to the planter's ambience.

Notes:

Notes:

Notes:

Notes:

Notes:

VLDB 2010 PLATINUM SPONSORS

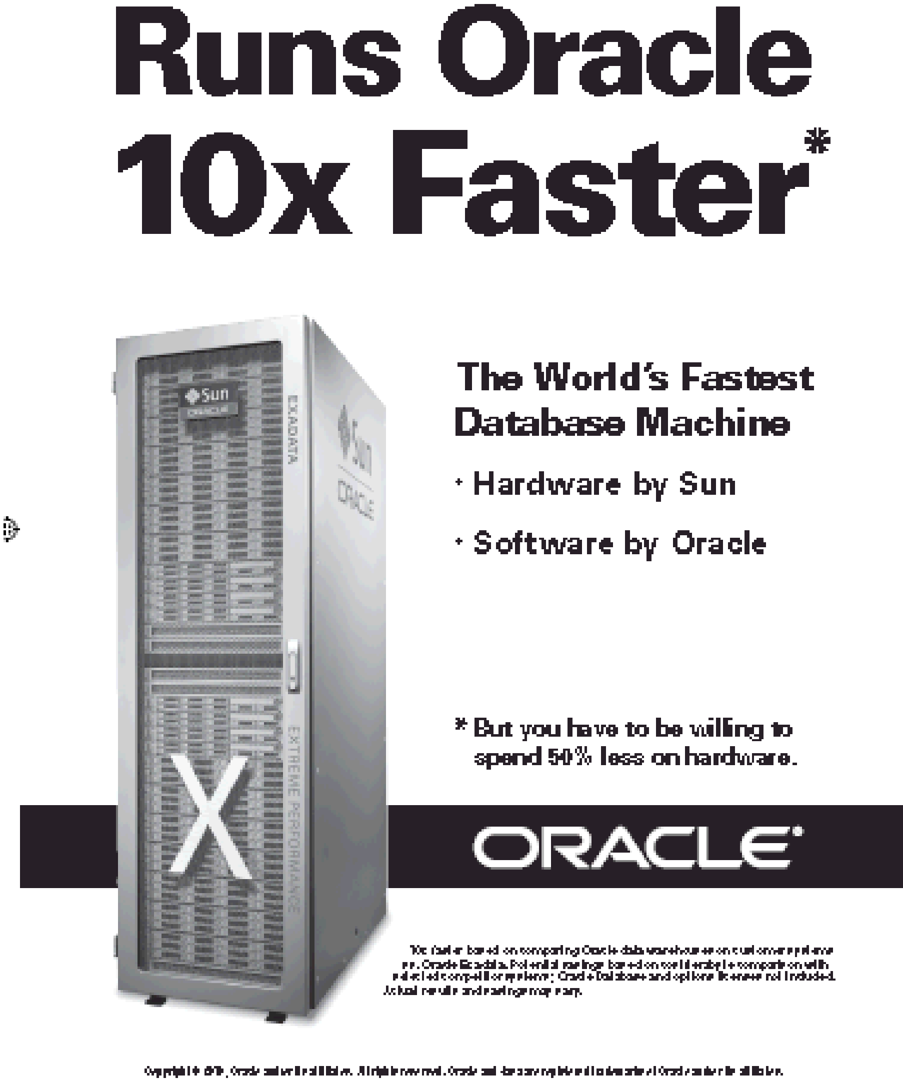
Runs Oracle
10x Faster*
The World's Fastest
Database Machine
• Hardware by Sun
• Software by Oracle
* But you have to be willing to
spend 50% less on hardware.
ORACLE
EXADATA
EXTREME PERFORMANCE

VLDB 2010 GOLD SPONSORS

IBM has a rich history in information management – relational databases were invented by IBM researchers in Silicon Valley. Since that time, IBM researchers have achieved many significant milestones in the field including key technologies resulting from projects such as System R, R*, Starburst, Garlic, DB2 pureXML, System S and Clio, as well as inventing the first data mining algorithms.

IBM Research is the leading corporate research institution in the world with a global community of 3,000 world-class scientists, whose rich history of discovery and innovation has brought international recognition. In addition to five Nobel prizes, IBM researchers have been recognized with six Turing awards, seven U.S. National Medals of Technology, five National Medals of Science and memberships in the National Academy of Sciences, and 11 inductions into National Inventors Hall of Fame. IBM Research has more than 59 members of the National Academy of Engineering and well over 300 industry organization fellows. IBM's own Frances Allen became the first woman to receive a Turing Award. For the seventeenth straight year, IBM received the most U.S. patents of any company worldwide.

IBM works with clients around the world to help their businesses grow "smarter" as the planet becomes more instrumented, interconnected and intelligent. Through its software, hardware, services, consulting and R&D expertise, IBM helps create systems that lead to less traffic, healthier food, cleaner water, and safer cities

About SAP

SAP is the world's leading provider of business software(*), offering applications and services that enable companies of all sizes and in more than 25 industries to become best-run businesses. With more than 97,000 customers in over 120 countries, the company is listed on several exchanges, including the Frankfurt stock exchange and NYSE, under the symbol "SAP." For more information, visit www.sap.com.

VLDB 2010 SILVER SPONSORS

VLDB 2010 DISCLAIMER

The organisers of VLDB 2010 have made every effort to ensure that the conference achieves its goal of disseminating the very best and most current information, advances and research on databases and information systems. Furthermore, the organisers have made every effort to ensure that the delegates remain comfortable and enjoy the experience of the conference. However, the organisers nor the VLDB Endowment take any responsibility for any damage, loss or inconvenience delegates may incur or experience in connection with the conference. In addition, the organisers cannot be held responsible for the correctness or appropriateness of the talks, papers, panels, tutorials and demonstrations included in the conference. In particular, changes to the published conference programme or cancellation of parts thereof do not entitle delegates to a full or partial refund of the conference fee.

Moreover, in the event of industrial disruption or other unforeseen circumstances, the organisers accept no responsibility for loss of monies incurred by delegates. The organisers accept no liability for injuries/losses of whatever nature incurred by participants and/or accompanying person, nor for loss or damage to their luggage and/or personal belongings. Delegates are expected to make their own arrangements with respect to personal insurance.

Any personal/business information supplied to the conference will be used by the conference organisation for the purposes of conference registration and administration. Names and addresses of delegates will be processed electronically and included in a list of delegates that may be posted and distributed during and in connection with the conference, unless the delegate has previously opted via the online registration system to not have their information shared. In addition, under control of the VLDB Endowment Board, this information may also be made available for purposes beneficial to the scientific community such as informing delegates of relevant publications, events, or items of interest to the community. By registering for the conference, delegates give their consent for such uses of their personal and business information.

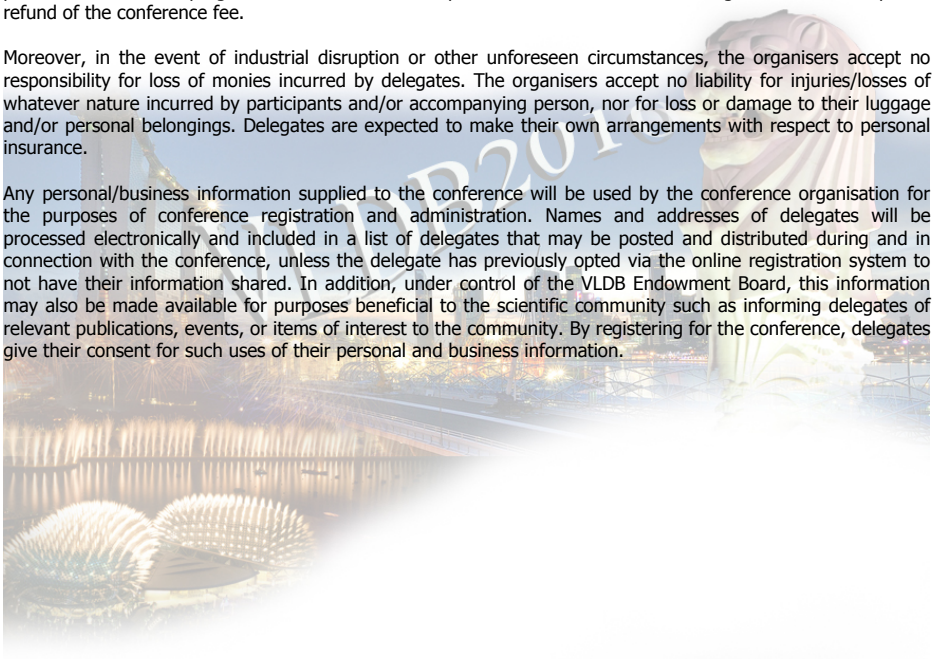

Platinum Sponsors

Microsoft ORACLE

Gold Sponsors

Google IBM SAP SYBASE

Silver Sponsors

ADSC YAHOO! hp TPC Transaction Processing Performance Council

Supporter

Sun microsystems

Supporting Organizations

Agency for Science, Technology and Research SINGAPORE

Singapore TOURISM BOARD

NUS National University of Singapore

NANYANG TECHNOLOGICAL UNIVERSITY

SMU SINGAPORE MANAGEMENT UNIVERSITY

Very Large Data Bases

Supporting Societies

1984

COMAD

DBSJ

CCF