SIGMOD/PODS 2011

June 12th-16th, 2011

Athens, Greece

Conference Program

ACM SIGMOD

Editors: *Tasos Kementsietsidis (IBM T.J. Watson), Yannis Velegrakis (University of Trento)*
Conference Proceedings Chairs

Template: *Yannis Velegrakis, Toronto 2004, Available for download at: http://db.disi.unitn.eu/pages/programBooklet*

2011 ACM SIGMOD/PODS CONFERENCE

June 12th – June 16th, 2011
Athens, Greece
www.sigmod2011.org

TABLE OF CONTENTS

Welcome message from the SIGMOD Chairs

Welcome to an exciting week in the city of Athens for the 2011 ACM SIGMOD Conference. Athens is a metropolitan and cosmopolitan city, with so many things to do and to see. It is also known as the birth place of Democracy, the city with the world-renown "Acropolis and Parthenon", with the famous Theater of Herodes Atticus and the "marble stadium" where the first modern time Olympic Games took place in 1896, home of Socrates, Plato, Pericles (Golden Age), and home of the very successful 2004 Olympic Games. And now the home of the 2011 ACM SIGMOD Conference! Athens is both an "ancient" and a "modern" city, in which visitors can walk safely and enjoy the rich - almost 5,000 year old - history it has to offer. The city offers a lot of sightseeing, museums, shopping and nightlife.

We have a program of several social events to complement an excellent technical program. The SIGMOD banquet is in the beautiful Island Restaurant by the sea, the SIGMOD reception on Monday at the Caravel Hotel (conference hotel) Roof Garden, and the PODS 30th Anniversary Colloquium and Reception on Sunday also at the hotel. Through the generous support from our sponsors (Platinum) EMC, Microsoft, and Oracle, (Gold) Google, IBM Research, SAP, Sybase, and Yahoo! Labs, (Silver) AsterData, HP, Intrasoft International, Kosmix, MarkLogic, Twitter, and VirtualTrip, and (Other Supporters) Greenplum and NEC, along with a contribution from ACM SIGMOD, we were able to keep the conference fees to a minimum with an extraordinarily low student registration that will allow many students to participate.

The scientific program was designed by over 150 dedicated volunteers. We would like to thank our chairs for all the time, knowledge, and insight they put into creating a very exciting program. In particular, we would like to thank our tutorial co-chairs Susan Davidson (University of Pennsylvania) and Tova Milo (Tel Aviv University), the keynote and panel chair Pat Selinger (IBM Research - Almaden), the industrial program co-chairs Richard Hull (IBM Research - Thomas J. Watson Research Ctr.) and Raghu Ramakrishnan (Yahoo! Research), the demonstration chair Chris Jermaine (Rice University), the workshop chair Christian S. Jensen (Aarhus University), and the undergraduate research program chair Irini Fundulaki (Institute of Computer Science - FORTH), along with our very wise advisory committee Laura Hass (IBM Research - Almaden) and Donald Kossmann (ETH Zurich). We also thank the SIGMOD Executive Committee for their guidance with special thanks to Lisa Singh for providing tremendously valuable history and advice.

We had 375 research papers submitted. A research program committee provided detailed reviews and extensive discussion following SIGMOD's double-blind reviewing policy. The program committee was led by nine group leaders: Sihem Amer-Yahia (Yahoo! Labs), Michael Böhlen (University of Zurich), Bettina Kemme (McGill University), Sam Madden (MIT), Jignesh Patel (University of Wisconsin), Dan Suciu (University of Washington), Wang-Chiew Tan (IBM Research - Almaden and UC Santa Cruz), Nesime Tatbul (ETH Zurich), and Min Wang (HP Labs China). The group leaders ensured that every paper had a champion and received thorough discussion. We accepted a record number of research papers. Nonetheless, it is clear there were still papers that were rejected that would have been valuable contributions to SIGMOD. The group leaders did a great job in keeping the discussions positive and reviewers focused on finding reasons to accept papers, rather than reasons to reject. We hope that SIGMOD continues the trend of accepting more papers to accommodate all the great ideas being produced by the community. All research papers have been invited to participate in SIGMOD's Experimental Repeatability effort, the goal of which is to help to enable SIGMOD papers to stand as reliable, archival work for future research.

A demonstration program committee of 35 people reviewed 81 proposals for system demonstrations and accepted 32. We will again be holding a Best Demonstration Award Competition to recognize the most innovative demonstrations. We encourage everyone to participate in the voting. An industrial program committee of 13 people reviewed 36 short presentation proposals, accepting 14. In addition, the industrial program will include two invited talks by Michael Abbott (Twitter) and Don Campbell (IBM). Rounding out the full program, we have a panel on data management issues in health and medical informatics, two invited talks by James Hamilton (Amazon) and Anastasia Ailamaki (EPFL), along with six tutorials on new applications of Datalog, flash memory, copy detection, data privacy, web data management, and statistical relational models.

SIGMOD continues its commitment to undergraduate research awarding eight scholarships to undergraduate researchers who will participate in the Undergraduate Research Poster session that will be co-located with a

Graduate Research Poster session. In addition, we will host the Third Annual SIGMOD Programming Contest. Student teams from degree granting institutions were invited to compete in this annual contest. This year, the task is to implement a high-throughput main-memory index that is made durable using a flash-based SSD.

On behalf of the conference organization committee, we would like to thank you all for your submissions to SIGMOD 2011, your hard work serving on SIGMOD 2011 committees, and for participating in the conference. It is your ideas, interest and input that shaped the conference and that continue to make it such a valuable venue for our community.

Timos Sellis	**Renée J. Miller**	**Anastasios Kementsietsidis**	**Yannis Velegrakis**
General Chair	*Program Chair*	*Proceedings Co-Chair*	*Proceedings Co-Chair*

Welcome Message from the PODS General Chair

It is our great pleasure to welcome you to the 2011 ACM Symposium on Principles of Database Systems – PODS'11. This year's symposium continues its tradition of being the premier international conference on the theoretical aspects of data management. PODS papers are distinguished by a rigorous approach to widely diverse problems in databases, often bringing to bear techniques from a variety of different areas, including computational logic, finite model theory, computational complexity, algorithm design and analysis, programming languages, and artificial intelligence. The first PODS conference was held in Los Angeles (CA) in 1982, with Jeffrey D. Ullman as General Chair. Since that time, virtually all new ideas, methods and techniques for data management have been investigated and presented in subsequent PODS conferences (see http://www09.sigmod.org/sigmod/pods/ for various information on the conference series).

To celebrate the 30th anniversary of the Symposium, the PODS Executive Committee has organized the PODS 30th Anniversary Colloquium, a special event held in June 12, 2011, with the goal of providing a retrospective on the role of database theory, and outlining a picture of the future directions of the discipline. The Colloquium featured five invited presentations from distinguished leaders in the field, namely:

- Moshe Y. Vardi: "The rise, fall, and rise of dependency theory: Part 1, the rise and fall",
- Ronald Fagin: "The rise, fall, and rise of dependency theory: Part 2, the rise from the ashes",
- Jeffrey D. Ullman: "Deductive Databases",
- Serge Abiteboul: "Trees, semistructured data, and other strange ways to go beyond tables",
- Victor Vianu: "Database Theory: Back to the Future".

The PODS Executive Committee is grateful to the speakers for their participation in the event, and to Frank Neven for organizing a lively discussion session ending the Colloquium.

As usual, putting together PODS'11 was a team effort. We are particularly grateful to the Program Chair, Thomas Schwentick, and to the whole program committee, who worked very hard in reviewing papers and providing feedback for authors. Finally, we thank the PODS proceedings and publicity Chair, Wim Martens, the hosting organization, and all our sponsors, in particular the ACM Special Interest Groups on Management of Data, for its invaluable support.

We hope that you will find the PODS'11 program interesting and thought provoking, in the best tradition of the PODS Symposium.

Maurizio Lenzerini
PODS General Chair

Welcome Message from the PODS PC Chair

This volume contains the proceedings of the Thirtieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2011), held in Athens, Greece, on June 13 – 15, 2011, in conjunction with the 2011 ACM SIGMOD International Conference on Management of Data.

The proceedings include a paper by Daniel Deutch and Tova Milo based on the keynote address by Tova Milo and two papers, the first by Marcelo Arenas and Jorge Pérez, based on the tutorial by Marcelo Arenas, and the second based on the tutorial by S. Muthu Muthukrishnan, and 25 contributed papers that were selected by the Program Committee from 113 submissions. Most of these papers are preliminary reports on work in progress. While they have been read by program committee members, they have not been formally refereed. Many of them will probably appear in more polished and detailed form in scientific journals.

The program committee selected the paper Data Exchange beyond Complete Data by Marcelo Arenas, Jorge Pérez and Juan L. Reutter for the PODS 2011 Best Paper Award. In addition, the announcement of the 2011 ACM PODS Alberto O. Mendelzon Test-of-Time Award appears in the proceedings. This year, the award is given to Optimal Aggregation Algorithms for Middleware by Ronald Fagin, Amnon Lotem, and Moni Naor. The paper originally appeared in the proceedings of PODS 2001. Warmest congratulations to the authors of these papers.

I thank all authors who submitted papers to the symposium, the members of the program committee and the many external referees for the enormous amount of work they have done. I am particularly grateful to Peter Widmayer for handling the papers for which I had a conflict of interest. The program committee did not meet in person, but carried out extensive discussions during the electronic PC meeting. I am grateful to Andrei Voronkov for his EasyChair system which made it so easy to manage and coordinate the discussion of the submitted papers.

I thank Maurizio Lenzerini, the PODS General Chair, for his constant help with so many details related to the organization of the symposium, to Wim Martens who served as the PODS Proceedings Chair and as PODS Publicity Chair, to Dirk van Gucht, the previous PODS Program Chair, and to Jan Paredaens, the previous PODS General Chair, for useful suggestions on several issues, to the members of the PODS Executive Committee, particularly to Jianwen Su as an important link between SIGMOD and PODS and to Renée J. Miller, the SIGMOD Program Chair, for the excellent collaboration.

I thank many colleagues involved in the organization of the conference for fruitful collaboration: Timos Sellis (SIGMOD General Chair); Lisa Singh (SIGMOD/PODS coordinator); Yannis Stavrakas (Web/Information Chair); Yannis Kotidis (Local Arrangements Chair) & Lisa Tolles (Sheridan Printing Company). Finally, I thank the SIGMOD/PODS sponsors for their support.

Thomas Schwentick
PODS Program Chair

CONFERENCE ORGANIZATION

General Chair:	Timos Sellis (IMIS/RC Athena)
Program Chair:	Renée J. Miller (University of Toronto)
Proceedings Chairs:	Anastasios Kementsietsidis (IBM Research – T.J. Watson Research Ctr.) Yannis Velegrakis (University of Trento)
Tutorials Chairs:	Susan Davidson (University of Pennsylvania) Tova Milo (Tel Aviv University)
Keynote/Panel Chair:	Pat Selinger (IBM Research - Almaden)
Industrial Program Chairs:	Richard Hull (IBM Research – T.J. Watson Research Ctr.) Raghu Ramakrishnan (Yahoo! Research)
Advisory Committee:	Laura Hass (IBM Research - Almaden) Donald Kossmann (ETH Zürich)
Demo Chair:	Chris Jermaine (Rice University)
New Researcher Symposium Chairs:	Yanlei Diao (University of Massachusetts) Christopher Ré (University of Wisconsin-Madison)
Workshop Chair:	Christian S. Jensen (Aarhus University)
Undergraduate Research Program Chair:	Irini Fundulaki (Institute of Computer Science - FORTH)
Finance Chair:	Mike Hatzopoulos (University of Athens)
Publicity Chairs:	Dieter Pfoser (Inst. for the Mgmt. of Info. Systems) Ariel Fuxman (Microsoft Research)
Sponsorship Chairs:	Mike Carey (University of California - Irvine) Vasilis Vassalos (Athens University of Economics & Business)
Exhibits Chairs:	Theodore Dalamagas (IMIS/RC Athena) Flip Korn (AT&T Labs - Research)
Local Arrangements Chairs:	Yannis Kotidis (Athens U. of Economics & Business) - Chair Anastasios Patrikakos (IMIS/RC Athena) Eva Paradissi (IMIS/RC Athena)
Registration Chair:	Alex Delis (University of Athens)
Local Demonstration Chair:	Maria Halkidi (University of Piraeus)
Web/Information Chair:	Yannis Stavrakas (IMIS/RC Athena)
Graduate Poster Chairs:	Stefan Schoenauer (University of Helsinki and HIIT) Mario A. Nascimento (University of Alberta) Dirk Van Gucht (Indiana University)
Program Committee Group Leaders:	Sihem Amer-Yahia (Yahoo! Labs) Michael Böhlen (University of Zurich) Bettina Kemme (McGill University) Sam Madden (MIT) Jignesh Patel (University of Wisconsin) Dan Suciu (University of Washington) Wang-Chiew Tan (IBM Research – Almaden & UC Santa Cruz) Nesime Tatbul (ETH Zurich) Min Wang (HP Labs China)

Program Committee:

Daniel Abadi (Yale University)
Ashraf Aboulnaga (University of Waterloo)
Sihem Amer-Yahia (Yahoo! Labs)
Periklis Andritsos (Ontario Cancer Institute)
Marcelo Arenas (PUC Chile)
Denilson Barbosa (University of Alberta)
Zohra Bellahsene (Université Montpellier)
Michael Böhlen (University of Zurich)
Angela Bonifati (CNR, Italy)
Malu Castellanos (HP)
Tiziana Catarci (Sapienza - Università di Roma)
Ugur Cetintemel (Brown University)
Yi Chen (Arizona State University)
Lei Chen (Hong Kong University of Science and Technology)
Hong Cheng (Chinese University of Hong Kong)
Mitch Cherniak (Brandeis University)
Laura Chiticariu (IBM Research - Almaden)
Panos Chrysanthis (University of Pittsburgh)
Bobby Cochrane (IBM)
Isabel Cruz (University of Illinois at Chicago)
Theodore Dalamagas (IMIS Institute/"Athena" R.C.)
Nilesh Dalvi (Yahoo! Research)
Umesh Dayal (HP)
Amol Deshpande (University of Maryland, College Park)
Yanlei Diao (University of Massachusetts, Amherst)
Gill Dobbie (University of Auckland)
Luna Dong (AT&T Labs - Research)
Alan Fekete (University of Sydney)
Enrico Franconi (Free University of Bolzano-Bozen)
Christoph Freytag (Humboldt-Universität zu Berlin)
Ada Fu (Chinese University of Hong Kong)
Irini Fundulaki (Institute of Computer Science - FORTH)
Ariel Fuxman (Microsoft Research)
Johann Gamper (Free University of Bolzano-Bozen)
Floris Geerts (University of Edinburgh)
Boris Glavic (University of Toronto)
Lukasz Golab (AT&T Labs - Research)
Jonathan Goldstein (Microsoft)
Dimitrios Gounopoulos (University of Athens)
Todd J. Green (University of California, Davis)
Melanie Herschel (University of Tubingen)
Wynne Hsu (National University of Singapore)
Ihab Ilyas (University of Waterloo)
Zack Ives (University of Pennsylvania)
Ryan Johnson (CMU)
Jaewoo Kang (Korea University)
Verene Kantere (Cyprus University of Technology)
Anastasios Kementsietsidis (IBM Research – T.J. Watson Research Ctr.)
Bettina Kemme (McGill University)
Maurice van Keulen (University of Twente)
Masura Kitsuregawa (The Univ. of Tokyo)
Solmaz Kolahi (University of British Columbia)
Georgia Koutrika (Stanford University)
Tim Kraska (UC Berkeley)
Alex Labrinidis (University of Pittsburgh)
Mong Li Lee (National University of Singapore)
Kristen Lefevre (University of Michigan)
Boon Thau Loo (University of Pennsylvania)
Qiong Luo (Hong Kong University of Science and Technology)
Sam Madden (MIT)
Volker Markl (University of Berlin)

Program Committee:

Gianni Mecca (Università della Basilicata)
Alexandra Meliou (University of Washington)
Sergey Melnik (Google)
Mirella Moro (Universidade Federal de Minas Gerais)
Felix Naumann (Hasso Plattner Institute)
Raymond Ng (University of British Columbia)
Joann Ordille (Avaya Labs)
Fatma Ozcan (IBM Research - Almaden)
Meral Ozsoyoglu (Case Western Reserve University)
Themis Palpanas (Univerity of Trento)
Paolo Papotti (Università Roma Tre)
Jignesh M. Patel (University of Wisconsin)
Neoklis Polyzotis (University of California, Santa Cruz)
Rachel Pottinger (University of British Columbia)
Ken Pu (UOIT)
Vibhor Rastogi (Yahoo! Research)
Flavio Rizzolo (Carleton University)
Elke Rundensteiner (Worcester Polytechnic Institute)
Ana Carolina Salgado (Federal University of Pernambucol)
Eric Simon (SAP)
Lisa Singh (Georgetown University)
Radu Sion (Stony Brook University)
Dan Suciu (University of Washington)
Wang-Chiew Tan (IBM Research - Almaden & UC Santa Cruz)
Nesime Tatbul (ETH Zurich)
Evimaria Terzi (Boston University)
Patrick Valduriez (INRIA, Montpellier)
Yannis Velegrakis (University of Trento)
Stratis Viglas (University of Edinburgh)
Min Wang (HP Labs China)
Kyu-Young Whang (Korea Adv. Inst. of Science & Technology)
Melanie Wu (Indiana University)
Xifeng Yan (University of California at Santa Barbara)
Calisto Zuzarte (IBM Canada)

Additional reviewers:

Azza Abouzied
Kamil Bajda-Pawlikowski
Spyros Blanas
Chiranjeeb Buragohain
Claudio Di Ciccio
Ricardo Ciferri
Oktie Hassanzadeh
Jiewen Huang
Norio Katayama
Peer Kröger
Frank McSherry
Tadashi Ohmori
Christopher Olston
Aditya Parameswaran
Antonella Poggi
Rajeev Rastogi
Kunihiko Sadakane
Monica Scannapieco
Alexander Thomson
Valeria Times
Steven Whang
Diego Zardetto

PODS Executive Committee:
Maurizio Lenzerini (Sapienza University of Rome, Italy)
Phokion G. Kolaitis (Univ. of California, Santa Cruz & IBM Research, USA)
Jan Paredaens (University of Antwerp, Belgium)
Thomas Schwentick (TU Dortmund University, Germany)
Jianwen Su (University of California, Santa Barbara, USA)
Dirk Van Gucht (Indiana University, USA)

PODS General Chair:
Maurizio Lenzerini (Sapienza University of Rome, Italy)

PODS Program Chair:
Thomas Schwentick (TU Dortmund University, Germany)

PODS Program Committee:
Sara Cohen (Hebrew University of Jerusalem, Israel)
Alin Deutsch (University of California, San Diego, USA)
Thomas Eiter (Vienna University of Technology, Austria)
Wenfei Fan (University of Edinburgh & Bell Labs, UK & USA)
Sudipto Guha (University of Pennsylvania, USA)
Claudio Gutierez (University of Chile, Chile)
Peter Haas (IBM Almaden Research Center, USA)
Ravi Kumar (Yahoo! Research, Silicon Valley, USA)
Domenico Lembo (University of Rome La Sapienza, Italy)
Thomas Lukasiewicz (Oxford University, UK)
Filip Murlak (Warsaw University, Poland)
Dan Olteanu (Oxford University, UK)
Rafail Ostrovsky (University of California Los Angeles, USA)
Thomas Schwentick (TU Dortmund University, Germany)
Cristina Sirangelo (ENS Cachan, France)
Stijn Vansummeren (Université Libre de Bruxelles, Belgium)
Gerhard Weikum (MPI Saarbrücken, Germany)
Peter Widmayer (ETH Zürich, Switzerland)
Ke Yi (Hong Kong University of Science and Technology, China)

PODS Proceedings and Publicity Chair:
Wim Martens (TU Dortmund University, Germany)

PODS External Reviewers:

Pankaj Agarwal
Shun'ichi Amano
Marcelo Arenas
Lars Arge
Diego Arroyuelo
Vince Bárány
Jérémy Barbay
Pablo Barceló
Tugkan Batu
Michael Benedikt
Leopoldo Bertossi
Vladimir Braverman
Gerth Stølting Brodal
Toon Calders
Andrea Calì
Diego Calvanese
James Cheney
Flavio Chierichetti
Rada Chirkova
Jan Chomicki
Francisco Claude
Nilesh Dalvi
Anirban Dasgupta
Claire David
Giuseppe De Giacomo
Daniel Deutch
Tommaso Di Noia
Yanlei Diao
Xiaoyong Du
Georges Dupret
Thomas Erlebach
Alexandre Evfimievski
Arash Farzan
Diego Figueira
Michael Fink
George H. L. Fletcher
Minos Garofalakis
Olivier Gauwin
Stephen Hegner
Joseph Hellerstein
André Hernich
Alexander Hinneburg
Hendrik Jan Hoogeboom
Tomas Hruz
Tomasz Idziaszek
Riko Jacob
Wojciech Jaworski
Benny Kimelfeld
Solmaz Kolahi
Leszek Kołodziejczyk
David Koop
Lukasz Kowalik
John Langford
Kasper Green Larsen
Silvio Lattanzi
Lap-Kei Lee
Dirk Leinders
Maurizio Lenzerini
Leonid Libkin
Anthony Widjaja Lin
Alejandro López-Ortiz
Maria Vanina Martinez
Frank McSherry
Massimo Mecella
Ulrich Meyer
Matúš Mihal'ák
Mohamed Mokbel
Jelani Nelson
Frank Neven
Evdokia Nikolova
Werner Nutt
Rodrigo Paredes
Paweł Parys
Jorge Pérez
Jeff Phillips
Reinhard Pichler
Lucian Popa
Gabriele Puppis
Leonardo Querzoni
Vibhor Rastogi
Juan Reutter
Riccardo Rosati
Sudeepa Roy
Michael Schmidt
Marcel Schöngens
Bernhard Seeger
Gerardo Simari
Yannis Sismanis
D. Sivakumar
Adam Smith
Eljas Soisalon-Soininen
Jacek Sroka
Heiner Stuckenschmidt
Tony Tan
Yufei Tao
Nikolaj Tatti
Bernhard Thalheim
Srikanta Tirthapura
David Toman
Denis Turdakov
Alejandro Vaisman
Sergei Vassilvitskii
Millist Vincent
Wei Wang
Jef Wijsen
David Woodruff
Norbert Zeh
Guido Zuccon

CONFERENCE BANQUET

The Conference Banquet will take place at the **Island**. It is located on the 27th km of Athens-Sounio Ave, 16672, Varkiza, Attika, Greece.

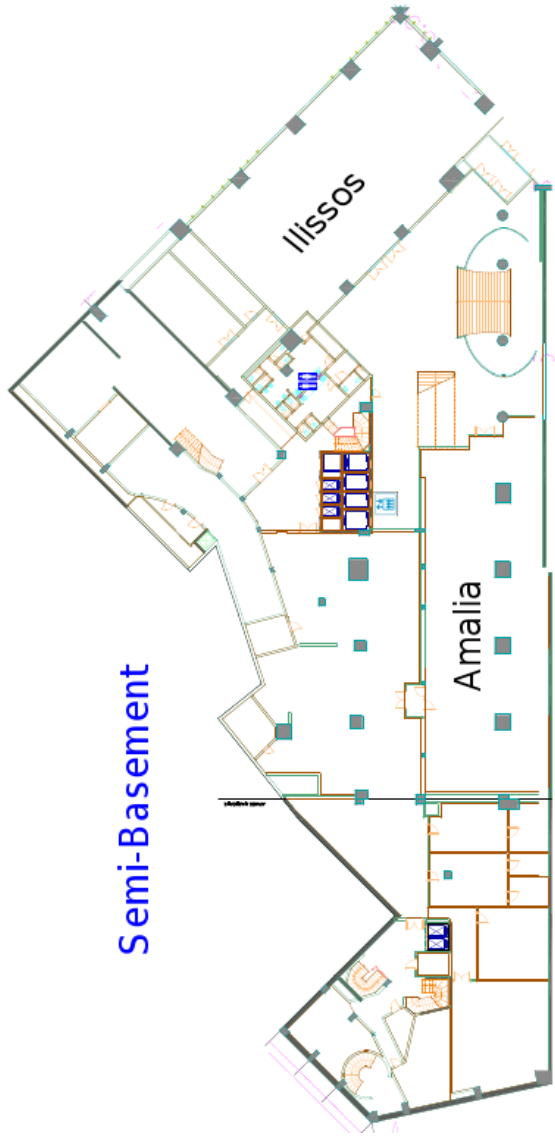
Semi-Basement
Amalia
Ilissos

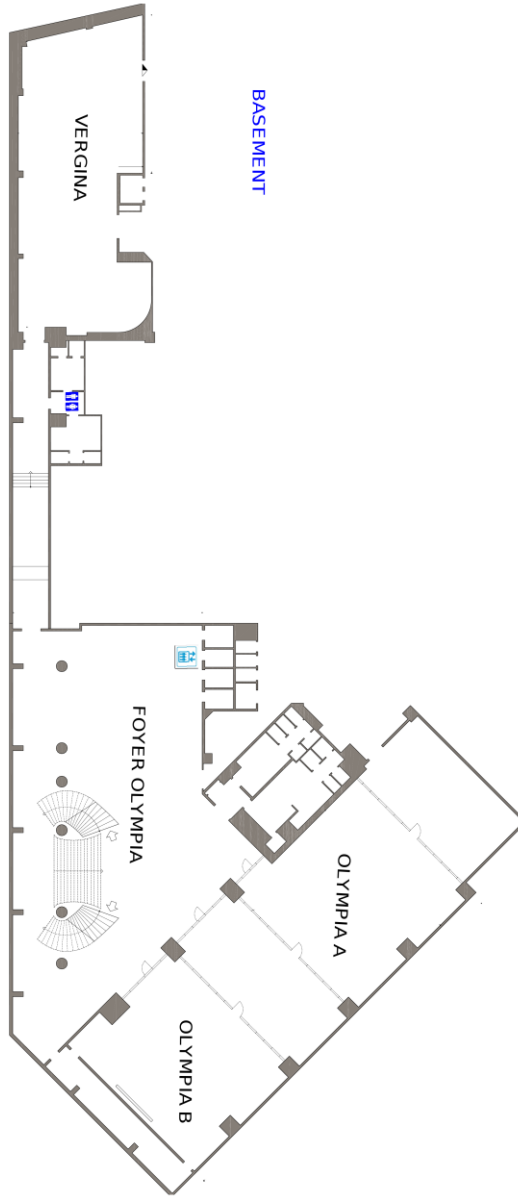
BASEMENT
VERGINA
FOYER OLYMPIA
OLYMPIA A
OLYMPIA B

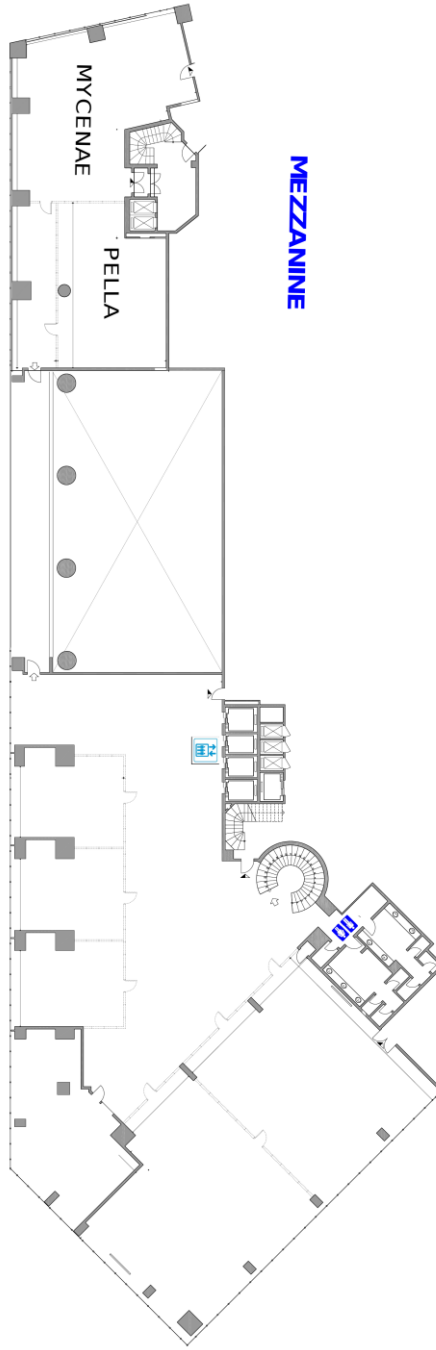
MEZZANINE
MYCENAE
PELLA

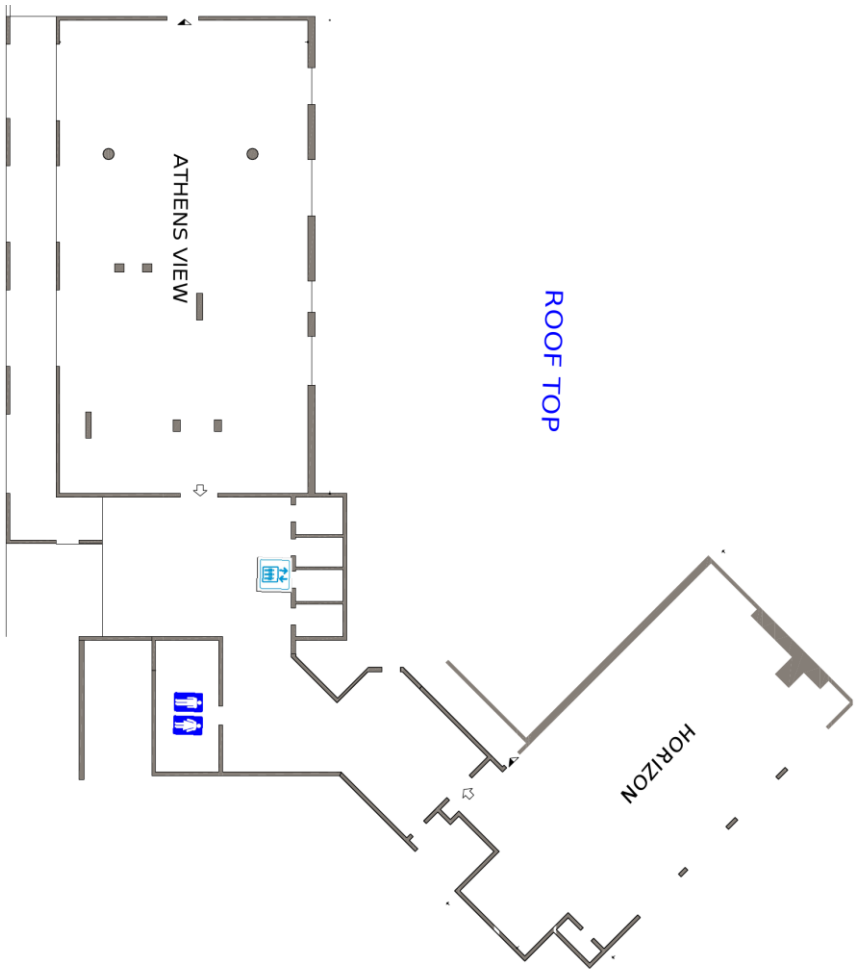
ATHENS VIEW
ROOF TOP
HORIZON

SUNDAY JUNE 12TH

	Olympia A	Olympia B	Vergina	Athens View	Pella	Horizon	Mycenae
08:00-09.00	**Continental Breakfast** (Foyer Olympia)						
09:00-10:30	**MobiDE Workshop**	**WebDB Workshop**	**NetDB Workshop**	NO SESSION	**DBSocial Workshop**	NO SESSION	**SWIM Workshop**
10:30-11:00	**Coffee Break** (Ballroom Foyer)						
11:00-12:30	**MobiDE Workshop**	**WebDB Workshop**	**NetDB Workshop**	NO SESSION	**DBSocial Workshop**	NO SESSION	**SWIM Workshop**
12:30-14:00	**Lunch Break** (Ilissos/Amalia) [provided for all the Workshops of the day participants]						
14:00-15:30	**MobiDE Workshop**	**WebDB Workshop**	**NetDB Workshop**	NO SESSION	**DBSocial Workshop**	NO SESSION	**SWIM Workshop**
15:30-16:00	**Coffee Break** (Ballroom Foyer)						
16:00-17:00	**MobiDE Workshop**	**WebDB Workshop**	**NetDB Workshop**	NO SESSION	**DBSocial Workshop**	NO SESSION	**SWIM Workshop**
17:30-20:00	**PODS 30th Anniversary Colloquium** (Horizon)						
20:00-22:00	**PODS Welcome Reception** (Roof Garden)						

MONDAY JUNE 13TH

	Olympia	Vergina	Athens View	Pella	Horizon	Mycenae
08:00-08.30	**Continental Breakfast** (Foyer Olympia)					
08:45-09:00	**PODS Opening** (Olympia)					
09:00-10:00	**PODS Keynote: : A Quest for Beauty and Wealth (or, Business Processes for Database Researchers)** Tova Milo, Tel Aviv University (Olympia)					
10:00-10:30	**Coffee Break** (Foyer Olympia)					
10:30-12:00	**PODS 1** Streaming and Sampling	NO SESSION	NO SESSION	**DaMoN Workshop**	NO SESSION	**DBTest Workshop**
12:00-13:30	**Lunch Break** (Ilissos/Amalia) [provided for the DaMoN and DBTest participants only]					
13:30-15:00	**PODS 2** Incomplete Information and Awards	NO SESSION	NO SESSION	**DaMoN Workshop**	NO SESSION	**DBTest Workshop**
15:00-15:30	**Coffee Break** (Foyer Olympia)					
15:30-17:00	**PODS 3** Index Structures and External Memory	NO SESSION	NO SESSION	**DaMoN Workshop**	NO SESSION	**DBTest Workshop**
17:00-17:30	**Coffee Break** (Foyer Olympia)					
17:30-19:00	NO SESSION	NO SESSION	NO SESSION	**DaMoN Workshop**	NO SESSION	**DBTest Workshop**
19:00-20:00	**Free time**					
20:00-22:30	**Welcome Reception, Undergraduate Poster Competition & Graduate Poster Session** (Roof Garden)					

TUESDAY JUNE 14TH

	Olympia A	Olympia B	Vergina	Athens View	Pella	Horizon	Mycenae
08:00-08.45	**Continental Breakfast** (Foyer Olympia)						
08:45-09:00	**Opening** (Olympia)						
09:00-10:00	**Keynote 1: Internet Scale Storage** James Hamilton, VP & Distinguished Engineer, Amazon Web Services (Olympia)						
10:00-10:30	**Coffee Break** (Foyer Olympia)						
10:30-11:30	**Research 1** Databases on New Hardware	**Research 2** Query Processing and Optimization	**Research 3** Schema Mapping and Data Integration	**Tutorial 1** Learning Statistical Models from Relational Data (Part 1)	**Industrial Keynote 1**	**PODS 4** Provenance	**Demo A** Ranking, the Web and Social Media
11:30-12:00					**Industrial 1** Data Management for Feeds and Streams		
12:00-13:30	**Lunch Break** (Ilissos/Amalia)						
13:30-15:00	**Research 4** Data on the Web	**Research 5** Data Privacy and Security	**Research 6** Data Consistency and Parallel DB	**Tutorial 1** Learning Statistical Models from Relational Data (Part 2)	**Industrial 2** Applying Hadoop	**PODS 5** Queries and Views	**Demo B** Systems and Performance **& Programming Contest**
15:00-15:30	**Coffee Break** (Foyer Olympia)						
15:30-16:30	**Research 7** Service Oriented Computing, Data Management in the Cloud	**Research 8** Spatial and Temporal Data Management	**Research 9** Shortest Paths and Sequence Data	**Tutorial 2** Datalog and Emerging Applications: An Interactive Tutorial	NO SESSION	**PODS Tutorial 1** Data Stream Analysis: Principles and Perspectives	**Demo C** Data Integration and Probabilistic Databases
16:30-17:30							
17:30-18:00	**Coffee Break** (Foyer Olympia)						
18:00-20:30	**New Researchers Symposium** (Olympia)						
20:30-23:00	**Microsoft Reception** (Foyer Olympia)						

WEDNESDAY JUNE 15TH

	Olympia A	Olympia B	Vergina	Athens View	Pella	Horizon	Mycenae
08:00-08.30	**Continental Breakfast** (Foyer Olympia)						
08:30-09:00	**Research 10** Data Provenance, Workflow and Cleaning	**Research 11** Information Extraction	**Research 12** Keyword Search and Ranked Queries	**Tutorial 3** Privacy-aware Data Management in Information Networks	**Industrial 3** Support for Business Analytics and Warehousing		**Demo D** User Support and Development Environments
09:00-10:00						**PODS Tutorial 2** Querying Semantic Web Data with SPARQL: State of the Art and Research Perspectives	
10:00-10:30	**Coffee Break** (Foyer Olympia)						
10:30-11:30	**Research 13** Stream and Complex Event Processing	**Research 14** Query Processing includes Best Paper	**Research 15** Data Mining	**Tutorial 4** Web Data Management	**Industrial Keynote 2**	**PODS 6** Semi-structured Data and XML	**Demo B** Systems and Performance **& Programming Contest**
11:30-12:00					**Industrial 4** Business Analytics		
12:00-13:30	**Lunch Break** (Ilissos/Amalia)						
13:30-14:00	**Bussiness Meeting** (Olympia)						
14:00-15:45	**Awards Presentations** (Olympia)						
15:45-16:00	**Coffee Break** (Foyer Olympia)						
16:00-17:00	**Plenary I** Posters of Papers in Research Sessions 11 - 21 (Ilissos)					**PODS 7** Rule-based Query Languages	**Demo A** Ranking, the Web and Social Media
17:00-18:00							
20:00-23:59	**Conference Banquet** (Island)						

THURSDAY JUNE 16TH

	Olympia A	Olympia B	Vergina	Athens View	Pella	Horizon	Mycenae
08:00-08.45	**Continental Breakfast** (Foyer Olympia)						
08:45-10.00	**Demo Awards and Keynote 2: Managing Scientific Data** Anastasia Ailamaki, EPFL (Olympia)						
10:00-10:30	**Coffee Break** (Ballroom Foyer)						
10:30-12:00	**Research 16** Information Retrieval	**Research 17** Probabilistic and Uncertain Databases	**Research 18** Scalable Data Analytics	**Tutorial 5** Data Management Over Flash Memory		**Panel 1** Data Management Issues in Health and Medical Informatics	**Demo D** User Support and Development Environments
12:00-13:30	**Lunch Break** (on your own)						
13:30-15:00	**Research 19** Graph Management	**Research 20** OLAP and its Applications	**Research 21** Similarity Search and Queries	**Tutorial 6** Large-Scale Copy Detection	**Industrial 5** Dynamic Optimization and Unstructured Content	NO SESSION	**Demo C** Data Integration and Probabilistic Databases
15:00-15:30	**Coffee Break** (Ballroom Foyer)						
15:30-17:30	**Plenary II** Posters of Papers in Research Sessions 1 - 10 (Ilissos)						

SESSION CONTENTS

SUNDAY 12TH 17:30 – 20:00

PODS 30th Anniversary Colloquium
Location: Horizon

SUNDAY 12TH 20:00 – 22:00

PODS Welcome Reception
Location: Roof Garden

MONDAY 13TH 08:45 – 10:00

PODS Opening and Keynote I
Location: Olympia
Session Chair: Maurizio Lenzerini

A Quest for Beauty and Wealth (or, Business Processes for Database Researchers)
Tova Milo (Tel Aviv University)

MONDAY 13TH 10:30 – 12:00

PODS Research 1: Streaming and Sampling
Location: Olympia
Session Chair: Ke Yi

Get the Most Out of Your Sample: Optimal Unbiased Estimators Using Partial Information
Edith Cohen (AT&T Labs-Research), Haim Kaplan (Tel Aviv University)

Tight Bounds for Lp Samplers, Finding Duplicates in Streams, and Related Problems
Hossein Jowhari, Mert Sağlam (Simon Fraser University), Gábor Tardos (Rényi Institute of Mathematics & Simon Fraser University)

Pan-private Algorithms Via Statistics on Sketches
Darakhshan Mir, S. Muthukrishnan, Aleksandar Nikolov, Rebecca N. Wright (Rutgers University)

FIFO Indexes for Decomposable Problems
Cheng Sheng, Yufei Tao (Chinese University of Hong Kong)

MONDAY 13TH 13:30 – 15:00

PODS Research 2: Incomplete Information and Awards
Location: Olympia
Session Chair: Victor Vianu

ACM PODS Alberto O. Mendelzon Test-of-Time Award: Optimal Aggregation Algorithms for Middleware.
Ronald Fagin, Amnon Lotem, and Moni Naor.
Award committee: Peter Buneman (University of Edinburgh), Meral Ozsoyoglu (Case Western Reserve University), Jianwen Su (UC Santa Barbara)

Best Paper Award: Data Exchange Beyond Complete Data
Marcelo Arenas (PUC Chile), Jorge Pérez (Universidad de Chile), Juan Reutter (University of Edinburgh)

Regular Paper: Incomplete Information and Certain Answers in General Data Models
Leonid Libkin (University of Edinburgh)

Regular Paper: Determining the Currency of Data
Wenfei Fan (University of Edinburgh & Harbin Institute of Technology), Floris Geerts (University of Edinburgh), Jef Wijsen (Université de Mons)

MONDAY 13TH 15:30 – 17:00

PODS Research 3: Index Structures and External Memory
Location: Olympia
Session Chair: Peter Widmayer

New Results on Two-dimensional Orthogonal Range Aggregation in External Memory
Cheng Sheng, Yufei Tao (Chinese University of Hong Kong)

On Finding Skylines in External Memory
Cheng Sheng , Yufei Tao (Chinese University of Hong Kong)

Beyond Simple Aggregates: Indexing for Summary Queries
Zhewei Wei, Ke Yi (Hong Kong University of Science and Technology)

Space-Efficient Substring Occurrence Estimation
Alessio Orlandi (University of Pisa), Rossano Venturini (ISTI-CNR)

MONDAY 13TH 20:00 – 22:30

SIGMOD Welcome Reception
Location: Roof Garden

Undergraduate Poster Competition
Location: Roof Garden

Study of techniques for the indirect enhancement of training data in the training of neural networks used in re-ranking of web search results
Panagiotis Parchas, School of Electrical and Computer Engineering, Department of Computer Science, National Technical University of Athens (NTUA), Greece

ATLAS: A Probabilistic Algorithm for High Dimensional Similarity Search
Jiaqi Zhai, Cornell University, USA

Outsourced Computation Verification
Roy Luo, University of California, Berkeley, USA

DYNAMO - Dynamic Web-Service Mashups
Apostolos Nydriotis, ECE Department, Technical University of Crete, Greece

DEMA: Dynamic Clustering of Spatio-Temporal Dataset to Improve Indexing Performance
Wook-Hee Kim, Electrical and Computer Engineering, Ulsan National Institute of Science and Technology, Korea

Link Prediction in Annotation Graphs using Graph Summarization
Philip Anderson, Computer Science Department, University of Maryland, USA

Graduate Posters Session
Location: Roof Garden

Research on Moving Objects with Multimodal Transportation Modes
Jianqiu Xu (FernUniversität Hagen)

Database Forensics in the Service of Information Accountability
Kyriacos E. Pavlou (The University of Arizona)

TUESDAY 14TH 08:45 – 10:00

SIGMOD Opening and Keynote I
Location: Olympia
Session Chair: Fatma Ozcan

Internet-Scale Storage
James Hamilton (Amazon Web Services)

TUESDAY 14TH 10:30 – 12:00

SIGMOD Research 1: Databases on New Hardware
Location: Olympia A
Session Chair: Yanlei Diao

LazyFTL: A Page-level Flash Translation Layer Optimized for NAND Flash Memory
Dongzhe Ma (Tsinghua University), Jianhua Feng (Tsinghua University), Guoliang Li (Tsinghua University)

Operation-Aware Buffer Management in Flash-based Systems
Yanfei Lv (Peking University); Bin Cui (Pku); Bingsheng He (Nanyang Technological Universi); Xuexuan Chen (Peking University)

SkimpyStash: RAM Space Skimpy Key-Value Store on Flash-based Storage
Biplob Debnath (EMC Corporation); Sudipta Sengupta (Microsoft Research); Jin Li (Microsoft Research)

Design and Evaluation of Main Memory Hash Join Algorithms for Multi-core CPUs
Spyros Blanas (Univ. of Wisconsin); Yinan Li (Univ. of Wisconsin); Jignesh Patel (Univ. of Wisconsin)

SIGMOD Research 2: Query Processing and Optimization
Location: Olympia B
Session Chair: Christoph Freytag

Query Optimization Techniques for Partitioned Tables
Herodotos Herodotou (Duke University); Nedyalko Borisov (Duke University); Shivnath Babu (Duke University)

CrowdDB: Answering Queries with Crowdsourcing
Michael Franklin (UC Berkeley); Donald Kossmann (ETH Zurich); Tim Kraska (UC Berkeley); Sukriti Ramesh (ETH Zurich); Reynold Xin (UC Berkeley)

Skyline Query Processing over Joins
Akrivi Vlachou (NTNU); Christos Doulkeridis (NTNU); Neoklis Polyzotis (UC - Santa Cruz)

Efficient Parallel Skyline Processing using Hyperplane Projections
Henning Koehler (University of Queensland); Jing Yang; Xiaofang Zhou (University of Queensland)

SIGMOD Research 3: Schema Mapping and Data Integration
Location: Vergina
Session Chair: Eric Simon

Automatic Discovery of Attributes in Relational Databases
Meihui Zhang (National University of Singapore); Marios Hadjieleftheriou (AT&T Labs - Research); Beng Chin Ooi (National University of Singapore); Cecilia Procopiuc (AT&T Labs - Research); Divesh Srivastava (AT&T Labs - Research)

Scalable Query Rewriting: A Graph-Based Approach
George Konstantinidis (ISI/USC); Jose Luis Ambite (ISI/USC)

Leveraging Query Logs for Schema Mapping Generation in U-MAP
Hazem Elmeleegy (Purdue University); Ahmed Elmagarmid (Qatar Computing Research Institute); Jaewoo Lee (Purdue University)

Designing and Refining Schema Mappings via Data Examples
Bogdan Alexe (UC Santa Cruz); Balder ten Cate (UC Santa Cruz); Phokion Kolaitis (UC Santa Cruz); Wang-Chiew Tan (IBM Research - Almaden and UC Santa Cruz)

SIGMOD Tutorial 1: Learning Statistical Models from Relational Data (Part 1)
Location: Athens View
Presenters: Lise Getoor (Univ. of Maryland); Lilyana Mihalkova (Univ. of Maryland)

SIGMOD Industrial Keynote and Session 1: Data Management for Feeds and Streams
Location: Pella
Session Chair: Raghu Ramakrishnan

SIGMOD Industrial Keynote 1
Michael Abbott (Twitter)

Bistro Data Feed Management System
Vladislav Shkapenyuk (AT&T Labs - Research), Theodore Johnson (AT&T Labs - Research), Divesh Srivastava (AT&T Labs - Research)

PODS Research 4: Provenance
Location: Horizon
Session Chair: Dan Olteanu

Provenance for Aggregate Queries
Yael Amsterdamer (Tel Aviv University and University of Pennsylvania), Daniel Deutch (Ben Gurion University and University of Pennsylvania), Val Tannen (University of Pennsylvania)

On Provenance Minimization
Yael Amsterdamer (Tel Aviv University and University of Pennsylvania), Daniel Deutch (Ben Gurion University and University of Pennsylvania), Tova Milo (Tel Aviv University), Val Tannen (University of Pennsylvania)

On the Complexity of Privacy-Preserving Complex Event Processing
Yeye He, Siddharth Barman (University of Wisconsin-Madison), Di Wang (Worcester Polytechnic Institute), Jeffrey F. Naughton (University of Wisconsin-Madison)

Provenance Views for Module Privacy
Susan B. Davidson, Sanjeev Khanna (University of Pennsylvania), Tova Milo (Tel Aviv University), Debmalya Panigrahi (Massachusetts Institute of Technology), Sudeepa Roy (University of Pennsylvania)

SIGMOD Demo A: Ranking, the Web and Social Media
Location: Mycenae

SkylineSearch: Semantic Ranking and Result Visualization for PubMed
Julia Stoyanovich (University of Pennsylvania), Mayur Lodha (Columbia University), William Mee (Columbia University), Kenneth A. Ross (Columbia University)

A Cross-Service Travel Engine for Trip Planning
Gang Chen (Zhejiang University), Chen Liu (National University of Singapore), Meiyu Lu (National University of Singapore), Beng Chin Ooi (National University of Singapore), Shanshan Ying (National University of Singapore), Anthony Tung (National University of Singapore), Dongxiang Zhang (National University of Singapore), Meihui Zhang (National University of Singapore)

BRRL: A Recovery Library for Main-Memory Applications in the Cloud
Tuan Cao (Cornell University), Benjamin Sowell (Cornell University), Marcos Vaz Salles (Cornell University), Alan Demers (Cornell University), Johannes Gehrke (Cornell University)

Tweets as Data: Demonstration of TweeQL and TwitInfo
Adam Marcus (MIT CSAIL), Michael S. Bernstein (MIT CSAIL), Osama Badar (MIT CSAIL), David R. Karger (MIT CSAIL), Samuel Madden (MIT CSAIL), Robert C. Miller (rcm@csail.mit.edu)

MOBIES: Mobile-Interface Enhancement Service for Hidden Web Database
Xin Jin (George Washington University), Aditya Mone (University of Texas at Arlington), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

Search Computing: Multi-domain Search on Ranked Data
Alessandro Bozzon (Politecnico di Milano), Daniele Braga (Politecnico di Milano), Marco Brambilla (Politecnico di Milano), Stefano Ceri (Politecnico di Milano), Francesco Corcoglioniti (Politecnico di Milano), Piero Fraternali (Politecnico di Milano), Salvatore Vadacca (Politecnico di Milano)

EnBlogue -- Emergent Topic Detection in Web 2.0 Streams
Foteini Alvanaki (Saarland University), Michel Sebastian (Saarland University), Krithi Ramamritham (IIT Bombay), Gerhard Weikum (Max-Planck Institute Informatics)

NOAM: News Outlets Analysis and Monitoring System
Ilias Flaounas (University of Bristol), Omar Ali (University of Bristol), Marco Turchi (European Commission), Tristan Snowsill (University of Bristol), Florent Nicart (Université de Rouen), Tijl De Bie (University of Bristol), Nello Cristianini (University of Bristol)

TUESDAY 14TH 13:30 – 15:00

SIGMOD Research 4: Data on the Web
Location: Olympia A
Session Chair: Luna Dong

Apples and Oranges: A Comparison of RDF Benchmarks and Real RDF Datasets
Songyun Duan (IBM Research - Thomas J. Watson Research Ctr.), Anastasios Kementsietsidis (IBM T.J. Watson), Kavitha Srinivas (IBM T. J. Watson), Octavian Udrea (IBM T. J. Watson)

Efficient Query Answering in Probabilistic RDF Graphs
Xiang Lian (HKUST), Lei Chen (Hong Kong University of Science and Technology)

Facet Discovery for Structured Web Search: A Query-log Mining Approach
Jeffrey Pound (University of Waterloo), Stelios Paparizos (Microsoft Research), Panayiotis Tsaparas

Schema-As-You-Go: On Probabilistic Tagging and Querying of Wide Tables
Meiyu Lu (NUS), Bing Tian Dai (NUS), Anthony Tung, Divyakant Agrawal (UC Santa Barbara)

SIGMOD Research 5: Data Privacy and Security
Location: Olympia B
Session Chair: Kristin LeFevre

No Free Lunch in Data Privacy
Daniel Kifer (Penn State), Ashwin Machanavajjhala (Yahoo)

TrustedDB: A Trusted Hardware based Database with Privacy and Data Confidentiality
Sumeet Bajaj (Stony Brook University), Radu Sion (Stony Brook University)

Differentially Private Data Cubes: Optimizing Noise Sources and Consistency
Bolin Ding (UIUC), Marianne Winslett (UIUC), Jiawei Han (UIUC); Zhenhui Li (UIUC)

iReduct: Differential Privacy with Reduced Relative Errors
Xiaokui Xiao (Nanyang Technological Univ), Gabriel Bender (Cornell University), Michael Hay (Cornell University), Johannes Gehrke (Cornell University)

SIGMOD Research 6: Data Consistency and Parallel DB
Location: Vergina
Session Chair: Tim Kraska

Fast Checkpoint Recovery Algorithms for Frequently Consistent Applications
Tuan Cao (Cornell University), Marcos Vaz Salles (Cornell University), Benjamin Sowell (Cornell University), Yao Yue (Cornell University), Johannes Gehrke (Cornell University), Alan Demers (Cornell University), Walker White (Cornell University)

A Latency and Fault-Tolerance Optimizer for Online Parallel Query Plans
Prasang Upadhyaya (University of Washington), YongChul Kwon (University of Washington), Magdalena Balazinska (University of Washington)

ArrayStore: A Storage Manager for Complex Parallel Array Processing
Emad Soroush (Univ. of Washington), Magdalena Balazinska (Univ. of Washington), Daniel Wang (SLAC)

Warding off the Dangers of Data Corruption with Amulet
Nedyalko Borisov (Duke University), Shivnath Babu (Duke University), Nagapramod Mandagere (IBM Almaden Research Center), Sandeep Uttamchandani (IBM Almaden Research Center)

SIGMOD Tutorial 1: Learning Statistical Models from Relational Data (Part 2)
Location: Athens View
Presenters: Lise Getoor (Univ. of Maryland); Lilyana Mihalkova (Univ. of Maryland)

SIGMOD Industrial 2: Applying Hadoop
Location: Pella
Session Chair: Nesime Tatbul

Apache Hadoop Goes Realtime at Facebook
Dhruba Borthakur (facebook), Jonathan Gray (facebook), Joydeep Sen Sarma (facebook), Kannan Muthukkaruppan (facebook), Nicolas Spiegelberg (facebook), Hairong Kuang (facebook), Karthik Ranganathan (facebook), Dmytro Molkov (facebook), Aravind Menon (facebook), Samuel Rash (facebook), Rodrigo Schmidt (facebook), Amitanand Aiyer (facebook)

Nova: Continuous Pig/Hadoop Workflows
Christopher Olston (Yahoo! Research), Greg Chiou (Yahoo!), Laukik Chitnis (Yahoo!), Francis Liu (Yahoo!), Yiping Han (Yahoo!), Mattias Larsson (Yahoo!), Andreas Neumann (Yahoo!), Vellanki B. N. Rao (Yahoo!), Vijayanand Sankarasubramanian (Yahoo!), Siddharth Seth (Yahoo!), Chao Tian (Yahoo!), Topher ZiCornell (Yahoo!), Xiaodan Wang (Johns Hopkins University)

A Hadoop Based Distributed Loading Approach to Parallel Data Warehouses
Yu Xu (Teradata), Pekka Kostamaa (Teradata), Yan Qi (Teradata), Jian Wen (UC Riverside), Kevin Keliang Zhao (UC San Diego)

A Batch of PNUTS: Experiences Connecting Cloud Batch and Serving Systems
Adam E Silberstein (Yahoo! Research), Russell Sears (Yahoo! Research), Wenchao Zhou (Univeristy of Pennsylvania), Brian Frank Cooper (Yahoo! Research)

PODS Research 5: Queries and Views
Location: Horizon
Session Chair: Cristina Sirangelo

Querying Graph Patterns
Pablo Barcelo (Universidad de Chile), Leonid Libkin, Juan L. Reutter (University of Edinburgh)

Maximizing Conjunctive Views in Deletion Propagation
Benny Kimelfeld, Jan Vondrák, Ryan Williams (IBM Research - Almaden)

Determining Relevance of Accesses at Runtime
Michael Benedikt, Georg Gottlob (University of Oxford), Pierre Senellart (Télécom ParisTech)

Parallel Evaluation of Conjunctive Queries
Paraschos Koutris, Dan Suciu (University of Washington)

SIGMOD Demo B: Systems and Performance & Programming Contest
Location: Mycenae

One-pass Data Mining Algorithms in a DBMS with UDFs
Carlos Ordonez (University of Houston), Sasi K Pitchaimalai (University of Houston)

Inspector Gadget: A Framework for Custom Monitoring and Debugging of Distributed Dataflows
Christopher Olston (Yahoo! Research), Benjamin Reed (Yahoo! Research)

RAFT at Work: Speeding-Up MapReduce Applications under Task and Node Failures
Jorge-Arnulfo Quiané-Ruiz (Saarland University), Christoph Pinkel (Saarland University), Jörg Schad (Saarland University), Jens Dittrich (Saarland University)

WattDB: An Energy-Proportional Cluster of Wimpy Nodes
Daniel Schall (TU Kaiserslautern), Volker Hudlet (TU Kaiserslautern)

WINACS: Construction and Analysis of Web-Based Computer Science Information Networks
Tim Weninger (University of Illinois Urbana-Champaign), Marina Danilevsky (University of Illinois Urbana-Champaign), Fabio Fumarola (Universita), Joshua Hailpern (University of Illinois Urbana-Champaign), Jiawei Han (University of Illinois Urbana-Champaign), Thomas J. Johnston (University of Illinois Urbana-Champaign), Surya Kallumadi (Kansas State University), Hyungsul Kim (University of Illinois Urbana-Champaign), Zhijin Li (University of Illinois Urbana-Champaign), David McCloskey (University of Illinois Urbana-Champaign), Yizhou Sun (University of Illinois Urbana-Champaign), Nathan E. TeGrotenhuis (Whitworth University), Chi Wang (University of Illinois Urbana-Champaign), Xiao Yu (University of Illinois Urbana-Champaign)

A Data-oriented Transaction Execution Engine and Supporting Tools
Ippokratis Pandis (Carnegie Mellon University), Pinar Tozun (Ecole Polytechnique Federale de Lausanne), Miguel Branco (Ecole Polytechnique Federale de Lausanne), Dimitris Karampinas (University of Patras), Danica Porobic (Ecole Polytechnique Federale de Lausanne), Ryan Johnson (University of Toronto), Anastasia Ailamaki (Ecole Polytechnique Federale de Lausanne)

iGraph in Action: Performance Analysis of Disk-Based Graph Indexing Techniques
Wook-Shin Han (Kyungpook National University), Minh-Duc Pham (Kyungpook National University), Jinsoo Lee (Kyungpook National University), Romans Kasperovics (Kyungpook National University), Jeffrey Xu Yu (Chinese University of Hong Kong)

StreamRec: A Real-Time Recommender System
Badrish Chandramouli (Microsoft Research), Justin J Levandoski (University of Minnesota), Ahmed Eldawy (University of Minnesota), Mohamed F Mokbel (University of Minnesota)

TUESDAY 14TH 15:30 – 17:30

SIGMOD Research 7: Service Oriented Computing, Data Management in the Cloud
Location: Olympia A
Session Chair: Ashraf Aboulnaga

Schedule Optimization for Data Processing Flows on the Cloud
Herald Kllapi (University Of Athens), Eva Sitaridi (University of Athens), Manolis Tsangaris (University of Athens), Yannis Ioannidis (University of Athens)

Zephyr: Live Migration in Shared Nothing Databases for Elastic Cloud Platforms
Aaron Elmore (UC Santa Barbara), Sudipto Das (UC Santa Barbara), Divyakant Agrawal (UC Santa Barbara), Amr El Abbadi (UC Santa Barbara)

Workload-Aware Database Monitoring and Consolidation
Carlo Curino (MIT), Evan Jones (MIT), Sam Madden (MIT), Hari Balakrishnan (MIT)

Predicting Cost Amortization for Query Services
Verena Kantere (Cyprus University of Technology), Debabrata Dash (ArcSight), Georgios Gratsias (ELCA Informatique SA), Anastasia Ailamaki (EPFL)

Performance Prediction for Concurrent Database Workloads
Jennie Duggan (Brown University), Ugur Cetintemel (Brown University), Olga Papaemmanouil (Brandeis University), Eli Upfal (Brown University)

SIGMOD Research 8: Spatial and Temporal Data Management
Location: Olympia B
Session Chair: Lei Chen

Reverse Spatial and Textual k Nearest Neighbor Search
Jiaheng Lu (Renmin University of China), Ying Lu (Renmin University of China), Gao Cong (Nanyang Technological University)

Location-Aware Type Ahead Search on Spatial Databases : Semantics and Efficiency
Senjuti Basu Roy (UT Arlington), Kaushik Chakrabarti (Microsoft Research)

Collective Spatial Keyword Querying
Xin Cao (Singapore NTU), Gao Cong (Nanyang Technological University), Christian Jensen (Aarhus University), Beng Chin Ooi (National University of Singapore)

Finding Semantics in Time Series
Peng Wang (Fudan University), Haixun Wang (Microsoft), Wei Wang (Fudan University)

Querying Contract Databases Based on Temporal Behavior
Elio Damaggio (UCSD); Alin Deutsch (UCSD); Dayou Zhou (UCSD)

SIGMOD Research 9: Shortest Paths and Sequence Data
Location: Vergina
Session Chair: Amelie Marian

Neighborhood-Privacy Protected Shortest Distance Computing in Cloud
Jun Gao (Peking University), Jeffrey Xu Yu (The Chinese University of Hong Kong), Ruoming Jin (Kent State University), Jiashuai Zhou, Tengjiao Wang, Dongqing Yang

On k-skip Shortest Paths
Yufei Tao (Chinese Univ. of Hong Kong), Cheng Sheng (The Chinese University of Hong Kong), JIan Pei (Simon Fraser Univ.)

Finding Shortest Path on Land Surface
Lian Liu (HKUST); Raymond Chi-Wing Wong (The Hong Kong University of Sc)

WHAM: A High-throughput Sequence Alignment Method
Yinan Li (Univ. of Wisconsin), Allison Terrell (Univ. of Wisconsin), Jignesh Patel (Univ. of Wisconsin)

A New Approach for Processing Ranked Subsequence Matching Based on Ranked Union
Wook-Shin Han (Kyungpook National University), Jinsoo Lee (Kyungpook National University), Yang-Sae Moon, Seung-won Hwang, Hwanjo Yu (POSTECH)

SIGMOD Tutorial 2: Datalog and Emerging Applications: An Interactive Tutorial
Location: Athens View
Presenters: Shan Shan Huang (LogicBlox), Todd J. Green (UC Davis), Boon Thau Loo (Univ. of Pennsylvania)

PODS Tutorial 1: Theory of Data Stream Computing: Where to Go
Location: Horizon (finishing at 16:30)
Presenter: S. Muthukrishnan (Rutgers University)
Session Chair: Dirk van Gucht

SIGMOD Demo C: Data Integration and Probabilistic Databases
Location: Mycenae

Pay-As-You-Go Mapping Selection in Dataspaces
Cornelia Hedeler (The University of Manchester), Khalid Belhajjame (The University of Manchester), Norman W Paton (The University of Manchester), Alvaro A A Fernandes (The University of Manchester), Suzanne M Embury (The University of Manchester), Lu Mao (The University of Manchester), Chenjuan Guo (The University of Manchester)

Exelixis: Evolving Ontology-Based Data Integration System
Haridimos Kondylakis (FORTH-ICS), Dimitris Plexousakis (FORTH-ICS)

U-MAP: A System for Usage-Based Schema Matching and Mapping
Hazem Elmeleegy (AT&T Labs - Research), Jaewoo Lee (Purdue University), El Kindi Rezig (Purdue University), Mourad Ouzzani (Purdue University), Ahmed Elmagarmid (Qatar Computing Research Institute, Qatar Foundation)

The SystemT IDE: An Integrated Development Environment for Information Extraction Rules
Laura Chiticariu (IBM Research - Almaden), Vivian Chu (IBM research - Almaden), Sajib Dasgupta (IBM Research - Almaden), Thilo W Goetz (IBM Software - Germany), Howard Ho (IBM Research - Almaden), Rajasekar Krishnamurthy (IBM Research - Almaden), Alexander Lang (IBM Software - Germany), Yunyao Li (IBM Research - Almaden), Bin Liu (University of Michigan), Sriram Raghavan (IBM Research - India), Frederick R Reiss (IBM Research - Almaden), Shivakumar Vaithyanathan (IBM Research - Almaden), Huaiyu Zhu (IBM Research - Almaden)

ProApproX: A Lightweight Approximation Query Processor over Probabilistic Trees
Pierre Senellart (Institut Télécom; Télécom ParisTech), Asma Souihli (Institut Télécom; Télécom ParisTech)

SPROUT²: A Squared Query Engine for Uncertain Web Data
Robert Fink (University of Oxford), Andrew Hogue (Google Inc.), Dan Olteanu (University of Oxford), Swaroop Rath (University of Oxford)

Fuzzy Prophet: Parameter exploration in uncertain enterprise scenarios
Oliver A Kennedy (EPFL), Steve Lee (Microsoft Corporation), Charles Loboz (Microsoft Corporation), Slawek Smyl (Microsoft Corporation), Suman Nath (Microsoft Research)

LinkDB: A Probabilistic Linkage Database System
Ekaterini Ioannou (Technical University of Crete), Wolfgang Nejdl (L3S Research Center), Claudia Niederée (L3S Research Center), Yannis Velegrakis (University of Trento)

TUESDAY 14TH 18:00 – 20:30

SIGMOD New Researchers Symposium
Location: Olympia

TUESDAY 14TH 20:30 – 23:00

Microsoft Reception
Location: Foyer Olympia

WEDNESDAY 15TH 08:30 – 10:00

SIGMOD Research 10: Data Provenance, Workflow and Cleaning
Location: Olympia A
Session Chair: Tasos Kementsietsidis

Interaction between Record Matching and Data Repairing
Wenfei Fan (University of Edinburgh), Jianzhong Li (Harbin Institute of Technology), Shuai Ma (Beihang University), Nan Tang (University of Edinburgh), Wenyuan Yu (University of Edinburgh)

We Challenge You to Certify Your Updates
Su Chen (National University of Singapore), Xin Luna Dong (AT&T Labs-Research), Laks V.S. Lakshmanan (University of British Columbia), Divesh Srivastava (AT&T Labs-Research)

Labeling Recursive Workflow Executions On-the-Fly
Zhuowei Bao (University of Pennsylvania), Susan B. Davidson (University of Pennsylvania), Tova Milo (Tel Aviv University)

Tracing Data Errors with View-Conditioned Causality
Alexandra Meliou (University of Washington), Wolfgang Gatterbauer (University of Washington), Suman Nath (Microsoft Research), Dan Suciu (University of Washington)

SIGMOD Research 11: Information Extraction
Location: Olympia B
Session Chair: Malu Castellanos

Hybrid In-Database Inference for Declarative Information Extraction
Daisy Zhe Wang (University of California, Berkeley), Michael J. Franklin (University of California, Berkeley), Minos Garofalakis (Technical University of Crete), Joseph M. Hellerstein (University of California, Berkeley), Michael L. Wick (University of Massachusetts, Amherst)

Faerie: Efficient Filtering Algorithms for Approximate Dictionary-based Entity Extraction
Guoliang Li (Tsinghua University), Dong Deng (Tsinghua University), Jianhua Feng (Tsinghua University)

Joint Unsupervised Structure Discovery and Information Extraction
Eli Cortez (Universidade Federal do Amazonas), Daniel Oliveira (Universidade Federal do Amazonas), Altigran S. da Silva (Universidade Federal do Amazonas), Edleno S. de Moura (Universidade Federal do Amazonas), Alberto H. F. Laender (Universidade Federal de Minas Gerais)

Attribute Domain Discovery for Hidden Web Databases
Xin Jin (George Washington University), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

SIGMOD Research 12: Keyword Search and Ranked Queries
Location: Vergina
Session Chair: Boris Glavic

Keyword Search over Relational Databases: A Metadata Approach
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), Raquel Trillo Lado (University of Zaragoza,), Yannis Velegrakis (University of Trento)

Sharing Work in Keyword Search over Databases
Marie Jacob (University Of Pennsylvania), Zachary Ives (University Of Pennsylvania)

Nearest Keyword Search in XML Documents
Yufei Tao (Chinese University of Hong Kong), Stavros Papadopoulos (Chinese University of Hong Kong), Cheng Sheng (Chinese University of Hong Kong), Kostas Stefanidis (Chinese University of Hong Kong)

Efficient and Generic Evaluation of Ranked Queries
Wen Jin (Independent Consultant), Jignesh M Patel (University of Wisconsin--Madison)

SIGMOD Industrial 3: Support for Business Analytics and Warehousing
Location: Pella
Session Chair: Yannis Papakonstantinou

Emerging Trends in the Enterprise Data Analytics
Fatma Özcan (IBM), David Hoa (IBM), Kevin S. Beyer (IBM), Andrey Balmin (IBM), Chuan Jie Liu (IBM), Yu Li (IBM)

Efficient Processing of Data Warehousing Queries in a Split Execution Environment
Kamil Bajda-Pawlikowski (Hadapt Inc. & Yale University), Daniel J Abadi (Hadapt Inc. & Yale University), Avi Silberschatz (Yale University), Erik Paulson (University of Wisconsin-Madison)

SQL Server Column Store Indexes
Per-Åke Larson (Microsoft), Cipri Clinciu (Microsoft), Eric N Hanson (Microsoft), Artem Oks (Microsoft), Susan L Price (Microsoft), Srikumar Rangarajan (Microsoft), Aleksandras Surna (Microsoft), Qingqing Zhou (Microsoft)

An Analytic Data Engine for Visualization in Tableau
Richard Wesley (Tableau Software), Matthew Eldridge (Tableau Software), Pawel T Terlecki (Tableau Software)

SIGMOD Tutorial 3: Privacy-aware Data Management in Information Networks
Location: Olympia A
Presenters: Michael Hay (Cornell University), Kun Liu (Yahoo! Labs), Gerome Miklau (University of Massachusetts Amherst), Jian Pei (Simon Fraser University), Evimaria Terzi (Boston University)

PODS Tutorial 2: Querying Semantic Web Data with SPARQL: State of the Art and Research Perspectives
Location: Horizon (starting at 09:00)
Session Chair: Alin Deutsch
Presenters: Marcelo Arenas (PUC Chile), Jorge Pérez (Universidad de Chile)

SIGMOD Demo D: User Support and Development Environments
Location: Mycenae

CONFLuEnCE: CONtinuous workFLow ExeCution Engine
Panayiotis Neophytou (University of Pittsburgh), Panos K. Chrysanthis (University of Pittsburgh), Alexandros Labrinidis (University of Pittsburgh)

Demonstration of Qurk: A Query Processor for HumanOperators
Adam Marcus (MIT CSAIL), Eugene Wu (MIT CSAIL), David R. Karger (MIT CSAIL), Samuel Madden (MIT CSAIL), Robert C. Miller (MIT CSAIL)

Automatic Example Queries for Ad Hoc Databases
Bill Howe (University of Washington), Garret Cole (University of Washington), Nodira Khoussainova (University of Washington), Leilani Battle (University of Washington)

NetTrails: A Declarative Platform for Maintaining and Querying Provenance in Distributed Systems
Wenchao Zhou (University of Pennsylvania), Qiong Fei (University of Pennsylvania), Shengzhi Sun (University of Pennsylvania), Tao Tao (University of Pennsylvania), Andreas Haeberlen (University of Pennsylvania), Zachary Ives (University of Pennsylvania), Boon Thau Loo (University of Pennsylvania), Micah Sherr (Georgetown University)

GBLENDER: Visual Subgraph Query Formulation Meets Query Processing
Changjiu Jin (Nanyang Technological University), Sourav S Bhowmick (Nanyang Technological University), Xiaokui Xiao (Nanyang Technological University), Byron Choi (Hong Kong Baptist University), Shuigeng Zhou (Fudan University)

Coordination through Querying in the Youtopia System
Nitin Gupta (Cornell University), Lucja Kot (Cornell University), Gabriel Bender (Cornell University), Sudip Roy (Cornell University), Johannes Gehrke (Cornell University), Christoph Koch (EPFL)

DBWiki: A Structured Wiki for Curated Data and Collaborative Data Management
Peter Buneman (University of Edinburgh), James Cheney (University of Edinburgh), Sam Lindley (University of Edinburgh), Heiko Müller (CSIRO)

Rapid Development of Web-Based Query Interfacesfor XML Datasets with QURSED
Abhijith Kashyap (SUNY at Buffalo), Michalis Petropoulos (SUNY at Buffalo)

WEDNESDAY 15TH 10:30 – 12:00

SIGMOD Research 13: Streams and Complex Event Processing
Location: Olympia A
Session Chair: Themis Palpanas

Changing Flights in Mid-air: A Model for Safely Modifying Continuous Queries
Kyumars Sheykh Esmaili (Systems Group), Tahmineh Sanamrad (ETH Zurich), Peter M Fischer (ETH Zurich), Nesime Tatbul (ETH Zurich)

How Soccer Players Would Do Stream Joins
Jens Teubner (ETH Zurich), Rene Mueller (IBM Almaden)

BE-Tree: An Index Structure to Efficiently Match Boolean Expressions over High-dimensional Discrete Space
Mohammad Sadoghi (University of Toronto), Hans-Arno Jacobsen (University of Toronto)

TI: An Efficient Indexing Mechanism for Real-Time Search on Tweets
Chun Chen (Zhejiang University), Feng Li (National University of Singapore), Beng Chin Ooi (National University of Singapore), Sai Wu (National University of Singapore)

SIGMOD Research 14: Query Processing
Location: Olympia B
Session Chair: Alexandra Meliou

More Efficient Datalog Queries: Subsumptive Tabling Beats Magic Sets
K. Tuncay Tekle (LogicBlox), Yanhong A. Liu (State University of New York at Stony Brook)

(Best Paper) Entangled Queries: Enabling Declarative Data-Driven Coordination
Nitin Gupta (Cornell University), Lucja Kot (Cornell University), Sudip Roy (Cornell University), Gabriel Bender (Cornell University), Johannes Gehrke (Cornell University), Christoph Koch (EPFL)

Data Generation using Declarative Constraints
Arvind Arasu (Microsoft Research), Raghav Kaushik (Microsoft Research), Jian Li (University of Maryland)

Efficient Auditing For Complex SQL queries
Raghav Kaushik (Microsoft Research), Ravi Ramamurthy (Microsoft Corporation)

SIGMOD Research 15: Data Mining
Location: Vergina
Session Chair: Julia Stoyanovich

Exact Indexing for Support Vector Machines
Hwanjo Yu (POSTECH (Pohang University of Science and Technology)), Ilhwan Ko (POSTECH (Pohang University of Science and Technology)), Youngdae Kim (POSTECH (Pohang University of Science and Technology)), Seungwon Hwang (POSTECH (Pohang University of Science and Technology)), Wook-Shin Han (Kyungpook National University)

Local Graph Sparsification for Scalable Clustering
Venu Satuluri (The Ohio State University), Srinivasan Parthasarathy (The Ohio State University), Yiye Ruan (The Ohio State University)

Advancing Data Clustering via Projective Clustering Ensembles
Francesco Gullo (University of Calabria), Carlotta Domeniconi (George Mason University), Andrea Tagarelli (University of Calabria)

Sampling Based Algorithms for Quantile Computation in Sensor Networks
Zengfeng Huang (Hong Kong University of Science and Technology), Lu Wang (Hong Kong University of Science and Technology), Ke Yi (Hong Kong University of Science and Technology), Yunhao Liu (Hong Kong University of Science and Technology)

SIGMOD Industrial Keynote and Session 4: Business Analytics
Location: Pella
Session Chair: Richard Hull

SIGMOD Industrial Keynote 2: Technology and Trends for Smarter Business Analytics
Don Campbell (IBM)

LCI: A Social Channel Analysis Platform for Live Customer Intelligence
Malu Castellanos (HP), Umeshwar Dayal (HP), Meichun Hsu (HP), Riddhiman Ghosh (HP), Mohamed Dekhil (HP), Yue Lu (UIUC), Lei Zhang (University of Illinois at Chicago)

SIGMOD Tutorial 4: Web Data Management
Location: Athens View
Presenters: Michael J Cafarella (University of Michigan), Alon Y Halevy (Google, Inc.)

PODS Research 6: Semistructured Data and XML
Location: Horizon (finishing at 11:40)
Session Chair: Thomas Schwentick

The Complexity of Text-Preserving XML Transformations
Timos Antonopoulos (Hasselt University and Transnational University of Limburg), Wim Martens (TU Dortmund), Frank Neven (Hasselt University and Transnational University of Limburg)

Efficient Evaluation for a Temporal Logic on Changing XML Documents
Mikolaj Bojańczyk (University of Warsaw), Diego Figueira (University of Warsaw & University of Edinburgh)

Finding a Minimal Tree Pattern Under Neighborhood Constraints
Benny Kimelfeld (IBM Research - Almaden), Yehoshua Sagiv (The Hebrew University)

SIGMOD Demo B: System and Performance & Programming Contest
Location: Mycenae

One-pass Data Mining Algorithms in a DBMS with UDFs
Carlos Ordonez (University of Houston), Sasi K Pitchaimalai (University of Houston)

Inspector Gadget: A Framework for Custom Monitoring and Debugging of Distributed Dataflows
Christopher Olston (Yahoo! Research), Benjamin Reed (Yahoo! Research)

RAFT at Work: Speeding-Up MapReduce Applications under Task and Node Failures
Jorge-Arnulfo Quiané-Ruiz (Saarland University), Christoph Pinkel (Saarland University), Jörg Schad (Saarland University), Jens Dittrich (Saarland University)

WattDB: An Energy-Proportional Cluster of Wimpy Nodes
Daniel Schall (TU Kaiserslautern), Volker Hudlet (TU Kaiserslautern)

WINACS: Construction and Analysis of Web-Based Computer Science Information Networks
Tim Weninger (University of Illinois Urbana-Champaign), Marina Danilevsky (University of Illinois Urbana-Champaign), Fabio Fumarola (Universita), Joshua Hailpern (University of Illinois Urbana-Champaign), Jiawei Han (University of Illinois Urbana-Champaign), Thomas J. Johnston (University of Illinois Urbana-Champaign), Surya Kallumadi (Kansas State University), Hyungsul Kim (University of Illinois Urbana-Champaign), Zhijin Li (University of Illinois Urbana-Champaign), David McCloskey (University of Illinois Urbana-Champaign), Yizhou Sun (University of Illinois Urbana-Champaign), Nathan E. TeGrotenhuis (Whitworth University), Chi Wang (University of Illinois Urbana-Champaign), Xiao Yu (University of Illinois Urbana-Champaign)

A Data-oriented Transaction Execution Engine and Supporting Tools
Ippokratis Pandis (Carnegie Mellon University), Pinar Tozun (Ecole Polytechnique Federale de Lausanne), Miguel Branco (Ecole Polytechnique Federale de Lausanne), Dimitris Karampinas (University of Patras), Danica Porobic (Ecole Polytechnique Federale de Lausanne), Ryan Johnson (University of Toronto), Anastasia Ailamaki (Ecole Polytechnique Federale de Lausanne)

iGraph in Action: Performance Analysis of Disk-Based Graph Indexing Techniques
Wook-Shin Han (Kyungpook National University), Minh-Duc Pham (Kyungpook National University), Jinsoo Lee (Kyungpook National University), Romans Kasperovics (Kyungpook National University), Jeffrey Xu Yu (Chinese University of Hong Kong)

StreamRec: A Real-Time Recommender System
Badrish Chandramouli (Microsoft Research), Justin J Levandoski (University of Minnesota), Ahmed Eldawy (University of Minnesota), Mohamed F Mokbel (University of Minnesota)

WEDNESDAY 15TH 13:30 – 14:00

SIGMOD Bussiness Meeting
Location: Olympia
Session Chair: Yannis Ioannidis

WEDNESDAY 15TH 14:00 – 15:45

SIGMOD Awards Presentation
Location: Olympia
Session Chair: Peter Buneman

WEDNESDAY 15TH 16:00 – 18:00

SIGMOD Plenary I: Posters of Research Papers from Sessions 11-21
Location: Ilissos

PODS Research 7: Rule-based Query Languages
Location: Horizon (finishing at 17:00)
Session Chair: Stijn Vansummeren

A Rule-based Language for Web Data Management
Serge Abiteboul (INRIA Saclay & LSV-ENS Cachan), Meghyn Bienvenu (CNRS & Université Paris Sud), Alban Galland, Émilien Antoine (INRIA Saclay & LSV-ENS Cachan)

Relational Transducers for Declarative Networking
Tom J. Ameloot, Frank Neven, Jan Van den Bussche (Hasselt University & Transnational University of Limburg)

Rewrite Rules for Search Database Systems
Ronald Fagin, Benny Kimelfeld, Yunyao Li (IBM Research - Almaden), Sriram Raghavan (IBM India Research Lab), Shivakumar Vaithyanathan (IBM Research - Almaden)

SIGMOD Demo A: Ranking, the Web and Social Media
Location: Mycenae

SkylineSearch: Semantic Ranking and Result Visualization for PubMed
Julia Stoyanovich (University of Pennsylvania), Mayur Lodha (Columbia University), William Mee (Columbia University), Kenneth A. Ross (Columbia University)

A Cross-Service Travel Engine for Trip Planning
Gang Chen (Zhejiang University), Chen Liu (National University of Singapore), Meiyu Lu (National University of Singapore), Beng Chin Ooi (National University of Singapore), Shanshan Ying (National University of Singapore), Anthony Tung (National University of Singapore), Dongxiang Zhang (National University of Singapore), Meihui Zhang (National University of Singapore)

BRRL: A Recovery Library for Main-Memory Applications in the Cloud
Tuan Cao (Cornell University), Benjamin Sowell (Cornell University), Marcos Vaz Salles (Cornell University), Alan Demers (Cornell University), Johannes Gehrke (Cornell University)

Tweets as Data: Demonstration of TweeQL and TwitInfo
Adam Marcus (MIT CSAIL), Michael S. Bernstein (MIT CSAIL), Osama Badar (MIT CSAIL), David R. Karger (MIT CSAIL), Samuel Madden (MIT CSAIL), Robert C. Miller (rcm@csail.mit.edu)

MOBIES: Mobile-Interface Enhancement Service for Hidden Web Database
Xin Jin (George Washington University), Aditya Mone (University of Texas at Arlington), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

Search Computing: Multi-domain Search on Ranked Data
Alessandro Bozzon (Politecnico di Milano), Daniele Braga (Politecnico di Milano), Marco Brambilla (Politecnico di Milano), Stefano Ceri (Politecnico di Milano), Francesco Corcoglioniti (Politecnico di Milano), Piero Fraternali (Politecnico di Milano), Salvatore Vadacca (Politecnico di Milano)

EnBlogue -- Emergent Topic Detection in Web 2.0 Streams
Foteini Alvanaki (Saarland University), Michel Sebastian (Saarland University), Krithi Ramamritham (IIT Bombay), Gerhard Weikum (Max-Planck Institute Informatics)

NOAM: News Outlets Analysis and Monitoring System
Ilias Flaounas (University of Bristol), Omar Ali (University of Bristol), Marco Turchi (European Commission), Tristan Snowsill (University of Bristol), Florent Nicart (Université de Rouen), Tijl De Bie (University of Bristol), Nello Cristianini (University of Bristol)

THURSDAY 16TH 08:45 – 10:00

SIGMOD Demo Awards and Keynote 2
Location: Olympia
Session Chair: Chris Jermaine

Managing Scientific Data
Anastasia Ailamaki (EPFL)

THURSDAY 16TH 10:30 – 12:00

SIGMOD Research 16: Information Retrieval
Location: Olympia A
Session Chair: Mirella Miro

Context-sensitive Ranking for Document Retrieval
Liang Jeff Chen (UC San Diego), Yannis Papakonstantinou (UC San Diego)

Score-Consistent Algebraic Optimization of Full-Text Search Queries with GRAFT
Nathan Bales (UC San Diego), Alin Deutsch (UC San Diego), Vasilis Vassalos (A.U.E.B)

Efficient Diversity-Aware Search
Albert Angel (University of Toronto), Nick Koudas (University of Toronto)

Mining a Search Engine's Corpus: Efficient Yet Unbiased Sampling and Aggregate Estimation
Mingyang Zhang (George Washington University), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

SIGMOD Research 17: Probabilistic and Uncertain Databases
Location: Olympia B
Session Chair: Yannis Kotidis

Ranking with Uncertain Scoring Functions: Semantics and Sensitivity Measures
Mohamed A Soliman (Greenplum), Ihab F Ilyas (University of Waterloo), Davide Martinenghi (Politecnico di Milano), Marco Tagliasacchi (Politecnico di Milano)

Querying Uncertain Data with Aggregate Constraints
Mohan Yang (Shanghai Jiao Tong University), Haixun Wang (Microsoft Research Asia), Haiquan Chen (Auburn University), Wei-Shinn Ku (Auburn University)

Jigsaw: Efficient optimization over uncertain enterprise data
Oliver A Kennedy (EPFL), Suman Nath (Microsoft Research)

Sensitivity Analysis and Explanations for Robust Query Evaluation in Probabilistic Databases
Bhargav Kanagal (University of Maryland), Jian Li (University of Maryland), Amol Deshpande (University of Maryland)

SIGMOD Research 18: Scalable Data Analytics
Location: Vergina
Session Chair: Daniel Abadi

Fast Personalized PageRank on MapReduce
Bahman Bahmani (Stanford University), Kaushik Chakrabarti (Microsoft Research), Dong Xin (Google Inc.)

Processing Theta-Joins using MapReduce
Alper Okcan (Northeastern University), Mirek Riedewald (Northeastern University)

Llama: Leveraging Columnar Storage for Scalable Join Processing in the MapReduce Framework
Yuting Lin (National University of Singapore), Divyakant Agrawal (University of California, Santa Barbara), Chun Chen (Zhejiang University), Beng Chin Ooi (National University of Singapore), Sai Wu (National University of Singapore)

A Platform for Scalable One-Pass Analytics using MapReduce
Boduo Li (University of Massachusetts Amherst), Edward Mazur (University of Massachusetts Amherst), Yanlei Diao (University of Massachusetts Amherst), Andrew McGregor (University of Massachusetts Amherst), Prashant Shenoy (University of Massachusetts Amherst)

SIGMOD Tutorial 5: Data Management over Flash Memory
Location: Athens View
Presenters: Ioannis Koltsidas (IBM Research), Stratis D Viglas (University of Edinburgh)

SIGMOD Panel 1: Data Management Issues in Health and Medical Informatics
Location: Horizon
Moderator: Tamer Özsu (University of Waterloo)
Panelists: Allen Brown (Microsoft), Helen Chen (Agfa Healthcare and University of Waterloo), Asuman Dogac (SRDC Ltd), Raymond Ng (UBC), Joel Saltz (Emory University), Wang Chiew Tan (IBM Research – Almaden and University of California – Santa-Cruz)

SIGMOD Demo D: User Support and Development Environments
Location: Mycenae

CONFLuEnCE: CONtinuous workFLow ExeCution Engine
Panayiotis Neophytou (University of Pittsburgh), Panos K. Chrysanthis (University of Pittsburgh), Alexandros Labrinidis (University of Pittsburgh)

Demonstration of Qurk: A Query Processor for HumanOperators
Adam Marcus (MIT CSAIL), Eugene Wu (MIT CSAIL), David R. Karger (MIT CSAIL), Samuel Madden (MIT CSAIL), Robert C. Miller (MIT CSAIL)

Automatic Example Queries for Ad Hoc Databases
Bill Howe (University of Washington), Garret Cole (University of Washington), Nodira Khoussainova (University of Washington), Leilani Battle (University of Washington)

NetTrails: A Declarative Platform for Maintaining and Querying Provenance in Distributed Systems
Wenchao Zhou (University of Pennsylvania), Qiong Fei (University of Pennsylvania), Shengzhi Sun (University of Pennsylvania), Tao Tao (University of Pennsylvania), Andreas Haeberlen (University of Pennsylvania), Zachary Ives (University of Pennsylvania), Boon Thau Loo (University of Pennsylvania), Micah Sherr (Georgetown University)

GBLENDER: Visual Subgraph Query Formulation Meets Query Processing
Changjiu Jin (Nanyang Technological University), Sourav S Bhowmick (Nanyang Technological University), Xiaokui Xiao (Nanyang Technological University), Byron Choi (Hong Kong Baptist University), Shuigeng Zhou (Fudan University)

Coordination through Querying in the Youtopia System
Nitin Gupta (Cornell University), Lucja Kot (Cornell University), Gabriel Bender (Cornell University), Sudip Roy (Cornell University), Johannes Gehrke (Cornell University), Christoph Koch (EPFL)

DBWiki: A Structured Wiki for Curated Data and Collaborative Data Management
Peter Buneman (University of Edinburgh), James Cheney (University of Edinburgh), Sam Lindley (University of Edinburgh), Heiko Müller (CSIRO)

Rapid Development of Web-Based Query Interfacesfor XML Datasets with QURSED
Abhijith Kashyap (SUNY at Buffalo), Michalis Petropoulos (SUNY at Buffalo)

THURSDAY 16TH 13:30 – 15:00

SIGMOD Research 19: Graph Management
Location: Olympia A
Session Chair: T.J. Green

Neighborhood Based Fast Graph Search in Large Networks
Arijit Khan (University of California, Santa Barbara), Nan Li (University of California, Santa Barbara), Xifeng Yan (University of California, Santa Barbara), Ziyu Guan (University of California, Santa Barbara), Supriyo Chakraborty (University of California, Los Angeles), Shu Tao (IBM T. J. Watson Research Center)

A Memory Efficient Reachability Data Structure Through Bit Vector Compression
Sebastiaan J. van Schaik (University of Oxford & Utrecht University), Oege de Moor (University of Oxford)

Incremental Graph Pattern Matching
Wenfei Fan (University of Edinburgh), Jianzhong Li (Harbin Institute of Technologies), Jizhou Luo (Harbin Institute of Technologies), Zijing Tan (Fudan University), Xin Wang (University of Edinburgh), Yinghui Wu (University of Edinburgh)

Assessing and Ranking Structural Correlations in Graphs
Ziyu Guan (University of California), Jian Wu (Zhejiang University), Qing Zhang (TaoBao.com), Ambuj Singh (University of California), Xifeng Yan (University of California)

SIGMOD Research 20: OLAP and its Applications
Location: Olympia B
Session Chair: Meral Ozsoyoglu

Graph Cube: On Warehousing and OLAP Multidimensional Networks
Peixiang Zhao (University of Illinois at Urbana-Champaign), Xiaolei Li (Microsoft Corporation), Dong Xin (Google Inc.), Jiawei Han (University of Illinois at Urbana-Champaign)

MaSM: Efficient Online Updates in Data Warehouses
Manos Athanassoulis (Ecole Polytechnique Federale de Lausanne), Shimin Chen (Intel Labs), Anastasia Ailamaki (Ecole Polytechnique Federale de Lausanne), Phillip B. Gibbons (Intel Labs), Radu Stoica (Ecole Polytechnique Federale de Lausanne)

Latent OLAP: Data Cubes over Latent Variables
Deepak Agarwal (Yahoo! Research), Bee-Chung Chen (Yahoo! Research)

E-Cube: Multi-Dimensional Event Sequence Analysis Using Hierarchical Pattern Query Sharing
Mo Liu (Worcester Polytechnic Institute), Elke Rundensteiner (Worcester Polytechnic Institute), Kara Greenfield (Worcester Polytechnic Institute), Chetan Gupta (Hewlett-Packard Labs), Song Wang (Hewlett-Packard Labs), Ismail Ari (Ozyegin University), Abhay Mehta (Hewlett-Packard Labs)

SIGMOD Research 21: Similarity Search and Queries
Location: Pella
Session Chair: Verena Kantera

ATLAS: A Probabilistic Algorithm for High Dimensional Similarity Search
Jiaqi Zhai (Cornell University), Yin Lou (Cornell University), Johannes Gehrke (Cornell University)

Flexible Aggregate Similarity Search
Yang Li (Shanghai Jiao Tong University), Feifei Li (Florida State University), Ke Yi (Hong Kong University of Science and Technology), Bin Yao (Florida State University), Min Wang (HP Labs China)

Effective Data Co-Reduction for Multimedia Similarity Search
Zi Huang (The University of Queensland), Hengtao Shen (The University of Queensland), Jiajun Liu (The University of Queensland), Xiaofang Zhou (The University of Queensland)

Efficient Exact Edit Similarity Query Processing with the Asymmetric Signature Scheme
Jianbin Qin (University of New South Wales), Wei Wang (University of New South Wales), Yifei Lu (University of New South Wales), Chuan Xiao (University of New South Wales), Xuemin Lin (University of New South Wales & East Normal China University)

SIGMOD Tutorial 6: Large-Scale Copy Detection
Location: Athens View
Presenters: Xin Luna Dong (AT&T Labs - Research), Divesh Srivastava (AT&T Labs - Research)

SIGMOD Industrial 5: Dynamic Optimization and Unstructured Content
Location: Pella
Session Chair: Russell Sears

Turbocharging DBMS Buffer Pool Using SSDs
Jaeyoung Do (University of Wisconsin-Madison), Donghui Zhang (Microsoft Jim Gray Systems Lab), Jignesh M Patel (University of Wisconsin-Madison), David J DeWitt (Microsoft Jim Gray Systems Lab), Jeffrey F Naughton (University of Wisconsin-Madison), Alan Halverson (Microsoft Jim Gray Systems Lab)

Online Reorganization in Read Optimized MMDBS
Felix Beier (Ilmenau University of Technology), Knut Stolze (IBM Research Development), Kai-Uwe Sattler (Ilmenau University of Technology)

Automated Partitioning Design in Parallel Database Systems
Rimma Nehme (Microsoft Jim Gray Systems Lab), Nicolas Bruno (Microsoft)

Oracle Database Filesystem
Krishna Kunchithapadam (Oracle Corporation), Wei Zhang (Oracle Corporation), Amit Ganesh (Oracle Corporation), Niloy Mukherjee (Oracle Corporation)

SIGMOD Demo C: Data Integration and Probabilistic Databases
Location: Mycenae

Pay-As-You-Go Mapping Selection in Dataspaces
Cornelia Hedeler (The University of Manchester), Khalid Belhajjame (The University of Manchester), Norman W Paton (The University of Manchester), Alvaro A A Fernandes (The University of Manchester), Suzanne M Embury (The University of Manchester), Lu Mao (The University of Manchester), Chenjuan Guo (The University of Manchester)

Exelixis: Evolving Ontology-Based Data Integration System
Haridimos Kondylakis (FORTH-ICS), Dimitris Plexousakis (FORTH-ICS)

U-MAP: A System for Usage-Based Schema Matching and Mapping
Hazem Elmeleegy (AT&T Labs - Research), Jaewoo Lee (Purdue University), El Kindi Rezig (Purdue University), Mourad Ouzzani (Purdue University), Ahmed Elmagarmid (Qatar Computing Research Institute, Qatar Foundation)

The SystemT IDE: An Integrated Development Environment for Information Extraction Rules
Laura Chiticariu (IBM Research - Almaden), Vivian Chu (IBM research - Almaden), Sajib Dasgupta (IBM Research - Almaden), Thilo W Goetz (IBM Software - Germany), Howard Ho (IBM Research - Almaden), Rajasekar Krishnamurthy (IBM Research - Almaden), Alexander Lang (IBM Software - Germany), Yunyao Li (IBM Research - Almaden), Bin Liu (University of Michigan), Sriram Raghavan (IBM Research - India), Frederick R Reiss (IBM Research - Almaden), Shivakumar Vaithyanathan (IBM Research - Almaden), Huaiyu Zhu (IBM Research - Almaden)

ProApproX: A Lightweight Approximation Query Processor over Probabilistic Trees
Pierre Senellart (Institut Télécom; Télécom ParisTech), Asma Souihli (Institut Télécom; Télécom ParisTech)

SPROUT²: A Squared Query Engine for Uncertain Web Data
Robert Fink (University of Oxford), Andrew Hogue (Google Inc.), Dan Olteanu (University of Oxford), Swaroop Rath (University of Oxford)

Fuzzy Prophet: Parameter exploration in uncertain enterprise scenarios
Oliver A Kennedy (EPFL), Steve Lee (Microsoft Corporation), Charles Loboz (Microsoft Corporation), Slawek Smyl (Microsoft Corporation), Suman Nath (Microsoft Research)

LinkDB: A Probabilistic Linkage Database System
Ekaterini Ioannou (Technical University of Crete), Wolfgang Nejdl (L3S Research Center), Claudia Niederée (L3S Research Center), Yannis Velegrakis (University of Trento)

THURSDAY 16TH 15:30 – 17:30

SIGMOD Plenary II: Posters of Research Papers from Sessions 1-10
Location: Ilissos

AWARDS

SIGMOD Contributions Award

Gerhard Weikum (Max-Planck-Institut für Informatik)

"For significant service to the database community, through research leadership and academic mentoring"

Gerhard Weikum is a Scientific Director at the Max Planck Institute for Informatics in Saarbruecken, Germany, where he is leading the department on databases and information systems. Earlier he held positions at Saarland University in Saarbruecken, Germany, at ETH Zurich, Switzerland, at MCC in Austin, Texas, and he was a visiting senior researcher at Microsoft Research in Redmond, Washington. He graduated from the University of Darmstadt, Germany. Gerhard Weikum's research spans transactional and distributed systems, self-tuning database systems, DB&IR integration, and automatic knowledge harvesting from Web and text sources. He co-authored a comprehensive textbook on transactional systems, received the VLDB 10-Year Award for his work on automatic DB tuning, and is one of the creators of the Yago knowledge base. Gerhard Weikum is an ACM Fellow, a Fellow of the German Computer Society, and a member of the German Academy of Science and Engineering. He has served on various editorial boards, including Communications of the ACM, and as program committee chair of conferences like ACM SIGMOD, Data Engineering, and CIDR. From 2003 through 2009 he was president of the VLDB Endowment.

SIGMOD Test-of-Time Award

Optimizing Queries Using Materialized Views: A Practical, Scalable Solution

Jonathan Goldstein (Microsoft Research) and Per-Åke Larson (Microsoft Research)

In this talk, we will discuss the evolution of materialized view technology in research and industry, beginning with the development of view maintenance algorithms, and query processing using views. In this context, we will discuss this year's test of time paper, which describes an efficient method for automatically exploiting a large class of materialized views to improve query performance. We will also outline the significant body of subsequent work on extending view maintenance and optimization support. The talk will conclude with a discussion of industry adoption of the work and the technical challenges in achieving further industry adoption.

"This paper presents a new technique for rewriting queries using materialized views that has influenced the design of query optimizers in at least two commercial database systems"

Jonathan Goldstein is a Principal Developer at Microsoft. His primary research interests are streaming, query processing, and optimization. During his 12 years at Microsoft, he has worked on query optimization (with Paul Larson), audio fingerprinting (with John Platt and Chris Burges), and cofounded (with Roger Barga) and then led the CEDR streaming research project, which became Microsoft StreamInsight, which is what he currently works on. Jonathan received his Ph.D. in 1999 from the University of Wisconsin.

Paul Larson is a Principal Researcher at Microsoft Research. His primary research interests are query optimization, query processing and design of database systems on modern hardware. Prior to joining Microsoft Research, he was a Professor in the Department of Computer Science at the University of Waterloo, Canada, for 15 years. Paul received his Ph.D. in 1976 from Åbo Akademi University in Finland. He is a Fellow of the ACM.

SIGMOD Best Paper Award

Entangled Queries: Enabling Declarative Data-Driven Coordination

Nitin Gupta (Cornell University), Lucja Kot (Cornell University), Sudip Roy (Cornell University), Gabriel Bender (Cornell University), Johannes Gehrke (Cornell University), Christoph Koch (EPFL)

Edgar F. Codd Innovations Award

Experiences with Problem #9

Surajit Chaudhuri (Microsoft Research)

In his Turing Award lecture, Jim Gray listed 12 information technology research problems. Problem #9 was to ensure that IT systems require minimal system administration overhead. In this talk, I will address Problem #9, as well as a few other problems that I have investigated over the years. I will explain how I became interested in these problems and what I learned from my exploration. The talk will also be an opportunity for me to thank my collaborators who played a critical role in my research projects over time. I will conclude by outlining a couple of problems that I still aspire to work on.

"For seminal contributions to research that led to practical tools for automated physical database design."

Surajit Chaudhuri is a Principal Researcher and the Research Area Manager overseeing data management research activities at Microsoft Research, Redmond. His areas of interest include self-tuning technology for databases, query optimization, data cleaning, and multi-tenant database systems. Working with his colleagues in Microsoft Research and the Microsoft SQL Server team, he helped incorporate the Index Tuning Wizard--and subsequently Database Engine Tuning Advisor--into Microsoft SQL Server. He initiated a project on data cleaning at Microsoft Research whose technology now ships in Microsoft SQL Server Integration Services. He also led the development of CMT, a conference management web service hosted by Microsoft Research since 1999 for the academic community. Surajit is an ACM Fellow, received the ACM SIGMOD Contributions Award in 2004 and received the VLDB 10-Year Best Paper Award in 2007. He is currently a member of the VLDB Endowment Board and the ACM SIGMOD Advisory Board. He was the Program Committee Chair for ACM SIGMOD 2006 and a Co-Chair of ACM SIGKDD 1999, and has served on the editorial boards of ACM TODS and IEEE TKDE. Surajit received his Ph.D. from Stanford University and B.Tech. from the Indian Institute of Technology, Kharagpur.

2011 SIGMOD Jim Gray Doctoral Dissertation Award

Database Cracking: Towards Auto-tuning Database Kernels.

Stratos Idreos (Centrum Wiskunde & Informatica)

Indices are heavily used in database systems in order to achieve the ultimate query processing performance. It takes a lot of time to create an index and the system needs to reserve extra storage space to store the auxiliary data structure. When updates arrive, there is also the overhead of maintaining the index. This way, which indices to create and when to create them has been and still is one of the most important research topics over the last decades. If the workload is known up-front or it can be predicted and if there is enough idle time to spare, then we can a priori create all necessary indices and exploit them when queries arrive. But what happens if we do not have this knowledge or idle time? Similarly, what happens if the workload changes often, suddenly and in an unpredictable way? Even if we can correctly analyze the current workload, it may well be that by the time we finish our analysis and create all necessary indices, the workload pattern has changed. Here we argue that a database system should just be given the data and queries in a declarative way and the system should internally take care of finding not only the proper algorithms and query plans but also the proper physical design to match the workload and application needs. The goal is to remove the role of database administrators, leading to systems that can completely automatically self-tune and adapt even to dynamic environments. Database Cracking implements the first adaptive kernel that automatically adapts to the access patterns by selectively and adaptively optimizing the data set purely for the workload at hand. It continuously reorganizes input data on-the-fly as a side-effect of query processing using queries as an advice of how data should be stored. Everything happens within operator calls during query processing and brings knowledge to the system that future operators in future queries can exploit. Essentially, the necessary indices are built incrementally as the system gains more and more knowledge about the workload needs.

"Idreos' dissertation proposes database cracking, a new paradigm for selecting and building indexes. Instead of deciding a-priori which indexes to build based on an anticipated query workload, the dissertation explores the idea of deferring such index construction until it becomes relevant for query processing, thus automating the difficult problem of physical database design. The dissertation is remarkable for its breadth, depth, and thorough implementation of the ideas in the open-source MonetDB database system."

Stratos holds a tenure track senior researcher position with CWI, the Dutch National Research Center for Mathematics and Computer Science. The main focus of his research is on adaptive query processing and database architectures, mainly in the context of column-stores. He also works on stream processing, distributed query processing and scientific databases. Stratos obtained his PhD from CWI and University of Amsterdam.In the past he has also been with the Technical University of Crete, Greece, and held research internship positions with Microsoft Research, Redmond, with EPFL, Switzerland and with IBM Almaden.

2011 SIGMOD Jim Gray Doctoral Dissertation Honorable Mention

Collaborative Data Sharing with Mappings and Provenance

Todd Green (University of Pennsylvania)

2011 SIGMOD Jim Gray Doctoral Dissertation Honorable Mention

On-line Index Selection for Physical Database Tuning

Karl Schnaitter (University of California - Santa Cruz)

PODS Best Paper Award

Data Exchange beyond Complete Data

Marcelo Arenas (PUC Chile), Jorge Pérez (Universidad de Chile), Juan L. Reutter (University of Edinburgh)

Marcelo Arenas is an Associate Professor at the Departmentof Computer Science at the Pontificia Universidad Catolica de Chile. He received his PhD from the University of Toronto in 2005. His research interests are in different aspects of database theory, such as expressive power of query languages, database semantics, inconsistency handling, database design, XML databases, data exchange, metadata management and database aspects of the Semantic Web. He has received an IBM Ph.D. Fellowship (2004), four best paper awards (PODS 2003, PODS 2005, ISWC 2006 and ICDT 2010) and an ACM-SIGMOD Dissertation Award Honorable Mention in 2006 for his Ph.D dissertation ``Design Principles for XML Data''. He has served on multiple program committees, and since 2009 he has been participating as an invited expert in the World Wide Web Consortium.

Jorge Pérez received a B.Sc. degree in Computer Engineering (2003), a M.Sc. degree in Computer Science (2004), and a Ph.D. degree in Computer Science (2011) from the P. Universidad Católica de Chile. He is currently affiliated with the Computer Science Department at Universidad de Chile. His research interests include data exchange and integration, schema mapping management, and the application of database tools and techniques to the Semantic Web. Pérez has received a Microsoft Research Ph.D. Fellowship (2009-2010), and three best paper awards (ISWC 2006, ESWC 2007, PODS 2011) for his work on database aspects of the Semantic Web, and on extensions of data exchange.

Juan Reutter received a Licenciate degree in Engineering Sciences (2007) and an M.Sc. degree in Computer Science (2009) from P. Universidad CatÃ³lica de Chile. He is currently a Ph.D. student at Edinburgh University under the supervision of Prof. Leonid Libkin. His research interest include data exchange and integration, schema mapping management, incomplete information and graph databases.

Alberto O. Mendelzon Test-of-Time PODS Award

Optimal Aggregation Algorithms for Middleware

Ronald Fagin (IBM Almaden Research Center), Amnon Lotem (University of Maryland-College Park), Moni Naor (Weizmann Institute of Science)

Assume that each object in a database has m grades, or scores, one for each of m attributes. For example, an object can have a color grade, that tells how red it is, and a shape grade, that tells how round it is. For each attribute, there is a sorted list, which lists each object and its grade under that attribute, sorted by grade (highest grade first). Each object is assigned an overall grade, that is obtained by combining the attribute grades using a fixed monotone "aggregation function", or "combining rule", such as min or average. To determine the top k objects, that is, k objects with the highest overall grades, the naive algorithm must access every object in the database, to find its grade under each attribute. Fagin has given an algorithm ("Fagin's Algorithm", or FA) that is much more efficient. For some monotone aggregation functions, FA is optimal with high probability in the worst case. We analyze an elegant and remarkably simple algorithm (the "threshold algorithm", or TA) that is optimal in a much stronger sense than FA. We show that TA is essentially optimal, not just for some monotone aggregation functions, but for all of them, and not just in a high-probability worst-case sense, but over every database. Unlike FA, which requires large buffers (whose size may grow unboundedly as the database size grows), TA requires only a small, constant-size buffer. TA allows early stopping, which yields, in a precise sense, an approximate version of the top k answers. We distinguish two types of access: sorted access (where the middleware system obtains the grade of an object in some sorted list by proceeding through the list sequentially from the top), and random access (where the middleware system requests the grade of object in a list, and obtains it in one step). We consider the scenarios where random access is either impossible, or expensive relative to sorted access, and provide algorithms that are essentially optimal for these cases as well.

"The paper investigates a very important problem that originates in multimedia databases: Given a set of objects with grades (rankings) on many attributes, find the objects with the best overall combined grades under some monotone combining function such as min or average. The paper presents a very simple algorithm, called Threshold Algorithm, and proves that it

is essentially optimal (in finding the best overall grades) for all monotone functions and over every database. Furthermore, the algorithm only requires a small constant-size buffer. The paper also gives adaptation of the algorithm for situations such as no random accesses. The Threshold Algorithm has been used in a wide range of applications where the problem naturally occurs, from databases with traditional and non-traditional types of data (music, video, text, uncertain data, etc.) to social networks, sensor networks, etc. The paper is among the most highly cited papers in PODS 2001, and perhaps all time. The paper has clearly had a major influence on the database and other research communities. Hence, the committee found it to be entirely worthy of this award."

Ronald Fagin is Manager of the Foundations of Computer Science group at the IBM Almaden Research Center, and is a member of the IBM Academy of Technology. He has won an IBM Corporate Award, eight IBM Outstanding Innovation Awards, an IBM Outstanding Technical Achievement Award, and two IBM key patent awards. He has published well over 100 papers, and has co-authored a book on "Reasoning about Knowledge". He has served on more than 30 conference program committees, including serving as Program Committee Chair of four different conferences. He received his B.A. in mathematics from Dartmouth College, and his Ph.D. in mathematics from the University of California at Berkeley. He was named a Fellow of IEEE for "contributions to finite-model theory and to relational database theory". He was named a Fellow of ACM for "creating the field of finite model theory and for fundamental research in relational database theory and in reasoning about knowledge". He was named a Fellow of AAAS (American Association for the Advancement of Science), for "fundamental contributions to computational complexity theory, database theory, and the theory of multi-agent systems". He was named Docteur Honoris Causa by the University of Paris, and a "Highly Cited Researcher" by ISI (the Institute for Scientific Information). He has won Best Paper awards at the 1985 International Joint Conference on Artificial Intelligence, the 2001 ACM Symposium on Principles of Database Systems, and the 2010 International Conference on Database Theory. He won a 2011 IEEE Technical Achievement Award "for pioneering contributions to the theory of rank and score aggregation". He was the winner of the 2004 ACM SIGMOD Edgar F. Codd Innovations Award, a lifetime achievement award in databases, for "fundamental contributions to database theory". In 2011, he won the IEEE Technical Achievement Award "for pioneering contributions to the theory of rank and score aggregation".

Amnon Lotem is the CTO of Skybox Security, a network security company based in San Jose, California. He has guided the company's research in cyber security solutions and is responsible for development of Skybox's patents, both awarded and pending. He has also led the company's development of core algorithms for attack simulation, risk analysis, and network compliance. His research interest is in AI planning, computer networks modeling, and cyber attack simulation. Amnon received his M.Sc. in Computer Science from Tel-Aviv University, and his Ph.D. in Computer Science from the University of Maryland, College Park. He has won Best Paper award at the 2001 ACM Symposium on Principles of Database Systems.

Moni Naor is an Israeli computer scientist, currently a professor at the Weizmann Institute of Science. Naor received his Ph.D. in 1989 at the University of California, Berkeley. His adviser was Manuel Blum. He works in various fields of computer science, mainly the foundations of cryptography. He is especially notable for creating non-malleable cryptography, Visual cryptography (with Adi Shamir), and suggesting various methods for verifying that users of a computer system are human (leading to the notion of CAPTCHA).

KEYNOTES

SIGMOD Keynote 1: Internet Scale Storage
James Hamilton (VP & Distinguished Engineer, Amazon Web Services)

The pace of innovation in data center design has been rapidly accelerating over the last five years, driven by the mega-service operators. I believe we have seen more infrastructure innovation in the last five years than we did in the previous fifteen. Most very large service operators have teams of experts focused on server design, data center power distribution and redundancy, mechanical designs, real estate acquisition, and network hardware and protocols. At low scale, with only a data or center or two, it would be crazy to have all these full time engineers and specialist focused on infrastructural improvements and expansion. But, at high scale with tens of data centers, it would be crazy not to invest deeply in advancing the state of the art.

Looking specifically at cloud services, the cost of the infrastructure is the difference between an unsuccessful cloud service and a profitable, self-sustaining business. With continued innovation driving down infrastructure costs, investment capital is available, services can be added and improved, and value can be passed on to customers through price reductions. Amazon Web Services, for example, has had eleven price reductions in four years. I don't recall that happening in my first twenty years working on enterprise software. It really is an exciting time in our industry.

I started working on database systems twenty years ago during a period of incredibly rapid change. We improved DB2 performance measured using TPC-A by a factor of ten in a single release. The next release, we made a further four-fold improvement. It's rare to be able to improve a product by forty fold in three years but, admittedly, one of the secrets is to begin from a position where work is truly needed. Back then, the database industry was in its infancy. Customers loved the products and were using them heavily, but we were not anywhere close to delivering on the full promise of the technology.

That's exactly where cloud computing is today--just where the database world was twenty years ago. Customers are getting great value from cloud computing but, at the same time, we have much more to do and many of the most interesting problems are yet to be solved. I could easily imagine tenfold improvement across several dimensions in over the next five years. What ties these two problems from different decades together is that some of the biggest problems in cloud computing are problems in persistent state management. What's different is that we now have to tackle these problems in a multi-tenant, high-scale, multi-datacenter environment. It's a new vista for database and storage problems.

In this talk, we'll analyze an internet-scale data center looking at the cost of power distribution, servers, storage, networking, and cooling on the belief that understanding what drives cost helps us focus on the most valuable research directions. We'll look at some of the fundamental technology limits approached in cloud database and storage solutions on the belief that, at scale, these limits will constrain practical solutions. And we'll consider existing cloud services since they form the foundation on which future solutions might be built.

James is VP and Distinguished Engineer at Amazon Web Services where he focuses on infrastructure efficiency, reliability, and scaling. Prior to AWS, James held leadership roles on several high-scale products and services including Windows Live, Exchange Hosted Services, Microsoft SQL Server, and IBM DB2. He loves all things server related and is interested in optimizing all components from data center power and cooling infrastructure, through server design, and the distributed software systems they host. James maintains high-scale services blog at http://perspectives.mvdirona.com.

SIGMOD Keynote 2: Managing Scientific Data
Anastasia Ailamaki (EPFL)

Today's scientific processes heavily depend on fast and accurate analysis of experimental data. Scientists are routinely overwhelmed by the effort needed to manage the volumes of data produced either by observing phenomena or by sophisticated simulations. As database systems have proven inefficient, inadequate, or insufficient to meet the needs of scientific applications, the scientific community typically uses special-purpose legacy software. When compared to a general-purpose DBMS, however, application-specific systems require more resources to maintain, and in order to achieve acceptable performance they often sacrifice data independence and hinder the reuse of knowledge. Nowadays, scientific datasets are growing at unprecedented rates, a result of increasing complexity of the simulated models and ever-improving instrument precision; consequently, scientistsqueries become more sophisticated as they try to interpret the data correctly. Datasets and scientific query complexity are likely to continue to grow indefinitely, rendering legacy systems increasingly inadequate. To respond to the challenge, the data management community aspires to solve scientific data management problems by carefully examining the problems of scientific applications and by developing special- or general-purpose scientific data management techniques and systems. This talk discusses the work of teams around the world in an effort to surface the most critical requirements of such an undertaking, and the technological innovations needed to satisfy them.

Anastasia (Natassa) Ailamaki is a Professor of Computer Sciences at the Ecole Polytechnique Federale de Lausanne (EPFL) in Switzerland. Her research interests are in database systems and applications, and in particular (a) in strengthening the interaction between the database software and the underlying hardware and I/O devices, including flash technology, and (b) in automating database design and computational database support for scientific applications. She has received a Finmeccanica endowed chair from the Computer Science Department at Carnegie Mellon (2007), a European Young Investigator Award from the European Science Foundation (2007), an Alfred P. Sloan Research Fellowship (2005), six best-paper awards at top conferences (2001-2006), and an NSF CAREER award (2002). She earned her Ph.D. in Computer Science from the University of Wisconsin-Madison in 2000.

SIGMOD Industrial Keynote 1:
Presenter: Michael Abbott (Twitter)

The last two years have seen a remarkable acceleration in online use of social networks, almost like the accelerating use of the web itself in the mid-nineties. This macro shift in online behavior presents new technical challenges in information search and discovery. Taking examples from data at Twitter, we will illustrate some of these challenges, our approach to them, and future directions during this session.

Mike Abbott (@mabb0tt) is Twitter's VP of engineering. He came to Twitter in 2010 from Palm, where he led the software development team that created HP/Palm's next-generation webOS platform. Abbott was previously the general manager for .NET online services at Microsoft. Prior to that, he co-founded Passenger Inc. and founded Composite Software. He holds a bachelor's degree from California Polytechnic State University, and has completed coursework toward a Ph.D. at the University of Washington.

SIGMOD Industrial Keynote 2: Technology and Trends for Smarter Business Analytics
Don Campbell (IBM)

Technology continues to change at an ever-increasing rate, and users continue to apply technologies in the workplace for business advantage. This whirlwind of change is having a major affect on business decision> making. This session will focus on trends and technologies impacting business intelligence and other analytics systems. We will explore the role of social networking, location intelligence, handheld devices, and other leading-edge technologies driving the future of smarter business decisions.

As CTO for IBM's Business Analytics division, Don Campbell oversees its research, technology, and innovation strategy. This includes social media, location intelligence, mobile computing, and information visualization. Don is a 23-year veteran of Cognos and IBM, and a member of IBM?s prestigious Academy of Technology.

PODS Keynote: A Quest for Beauty and Wealth (or, Business Processes for Database Researchers)

Tova Milo (Tel Aviv University)

While classic data management focuses on the data itself, research on Business Processes considers also the context in which this data is generated and manipulated, namely the processes, the users, and the goals that this data serves. This allows the analysts a better perspective of the organizational needs centered around the data. As such, this research is of fundamental importance.

Much of the success of database systems in the last decade is due to the beauty and elegance of the relational model and its declarative query languages, combined with a rich spectrum of underlying evaluation and optimization techniques, and efficient implementations. This, in turn, has lead to an economic wealth for both the users and vendors of database systems. Similar beauty and wealth are sought for in the context of Business Processes. Much like the case for traditional database research, elegant modeling and rich underlying technology are likely to bring economic wealth for the Business Process owners and their users; both can benefit from easy formulation and analysis of the processes. While there have been many important advances in this research in recent years, there is still much to be desired: specifically, there have been many works that focus on the processes behavior (flow), and many that focus on its data, but only very few works have dealt with both. We will discuss here the important advantages of a holistic flow-and-data framework for Business Processes, the progress towards such a framework, and highlight the current gaps and research directions.

Tova Milo received her Ph.D. degrees in Computer Science from theHebrew University, Jerusalem, in 1992. After graduating she workedat the INRIA research institute in Paris and at University ofToronto and returned to Israel in 1995, joining the School ofComputer Science at Tel Aviv university where she is now a full Professor. Her research focuses on advanced database applicationssuch as data integration, XML and semi-structured information,Web-based applications and Business Processes, studying both theoretical and practical aspects. Tova served as the Program Chairof several international conferences, including PODS, ICDT, VLDB, XSym, and WebDB. She is a member of the VLDB Endowment and the ICDT executive board and is an editor of TODS, the VLDB Journal and the Logical Methods in Computer Science Journal. She has received grants from the Israel Science Foundation, the US-Israel Binational Science Foundation, the Israeli and French Ministry of Science and theEuropean Union, and is a recipient of the 2010 ACM PODS Alberto O.Mendelzon Test-of-Time Award.

TUTORIALS

SIGMOD Tutorial 1: Learning Statistical Models from Relational Data

Location: Athens View

Presenters: Lise Getoor (Univ. of Maryland); Lilyana Mihalkova (Univ. of Maryland)

Statistical Relational Learning (SRL) is a subarea of machine learning which combines elements from statistical and probabilistic modeling with languages which support structured data representations. In this survey, we will: 1) provide an introduction to SRL, 2) describe some of the distinguishing characteristics of SRL systems, including relational feature construction and collective classification, 3) describe three SRL systems in detail, 4) discuss applications of SRL techniques to important data management problems such as entity resolution, selectivity estimation, and information integration, and 5) discuss connections between SRL methods and existing database research such as probabilistic databases.

Lise Getoor is an Associate Professor in the Computer Science Department at the University of Maryland, College Park. Her main research areas are machine learning and reasoning under uncertainty and she has also done work in database management, social network analysis and visual analytics. She is PC co-chair of ICML 2011, and has served as senior PC or PC member for conferences including AAAI, ICML, IJCAI, ICWSM, KDD, SIGMOD, UAI, VLDB, and WWW. She is an ACM Transactions on Knowledge Discovery and Data Associate Editor, was a Journal of Artificial Intelligence Research Associate Editor, and Machine Learning Journal Action Editor. She is on the board of the International Machine Learning Society, and has served on the AAAI Council. She is a recipient of an NSF Career Award, was a Microsoft New Faculty Fellow finalist and was awarded a National Physical Sciences Consortium Fellowship. She received her PhD from Stanford University, her MS from University of California, Berkeley, and her BS from the University of California, Santa Barbara.

Lilyana Mihalkova is a "Computing Innovations" post-doctoral fellow at the University of Maryland, College Park. She received her PhD from the University of Texas at Austin. Her research interests are in artificial intelligence and machine learning, including statistical relational learning, transfer learning, reasoning under uncertainty, and applications to Web and social media domains.

SIGMOD Tutorial 2: Datalog and Emerging Applications: An Interactive Tutorial

Location: Athens View

Presenters: Shan Shan Huang (LogicBlox), Todd J. Green (UC Davis), Boon Thau Loo (Univ. of Pennsylvania)

We are witnessing an exciting revival of interest in recursive Datalog queries in a variety of emerging application domains such as data integration, information extraction, networking, program analysis, security, and cloud computing. This tutorial briefly reviews the Datalog language and recursive query processing and optimization techniques, then discusses applications of Datalog in three application domains: data integration, declarative networking, and program analysis. Throughout the tutorial, we use LogicBlox, a commercial Datalog engine for enterprise software systems, to allow the audience to walk through code examples presented in the tutorial.

Shan Shan Huang is the lead of language and compiler technologies at LogicBlox, Inc. She oversees the design and development of DatalogLB -- a dialect of Datalog with extensions to support the expressiveness, modularity, and reuse necessary for developing enterprise decision support applications. These extensions include integrity constraints, module system, generic programming facilities, view-update support, and state. Shan Shan received her B.S. in Computer Science from Massachusetts Institute of Technology in 2000, and her Ph.D. in Computer Science from Georgia Institute of Technology in 2008. During her Ph.D., she was the recipient of the NSF Ph.D. Fellowship and the Intel Ph.D. Fellowship. Prior to her Ph.D. studies, Shan Shan worked for three years in software development at Fidelity Investments and ArsDigita, Inc.

T.J. Green is an Assistant Professor of Computer Science in the Department of Computer Science at the University of California, Davis. His research interests include collaborative data sharing, data provenance, incomplete and probabilistic databases, query optimization, semistructured data, and streaming data processing. He received his B.S. in Computer Science from Yale University in 1997, his M.S. in Computer Science from the University of Washington in 2001, and his Ph.D. in Computer and Information Science from the University of Pennsylvania in 2009. As a PhD student, he won the Best Student Paper award at ICDT 2009, and the Morris and Dorothy Rubinoff Award in 2010 (awarded to the outstanding computer science dissertation from the University of Pennsylvania). Prior to beginning his Ph.D., he worked for four years at Microsoft as a Software Design Engineer and Development Lead, and for two years at Xyleme as a Software Design Engineer.

Boon Thau Loo is an Assistant Professor in the Computer and Information Science department at the University of Pennsylvania. He received his Ph.D. degree in Computer Science from the University of California at Berkeley in 2006. Prior to his Ph.D, he received his M.S. degree from Stanford University in 2000, and his B.S. degree with highest honors from UC Berkeley in 1999. His research focuses on distributed data management systems, Internet-scale query processing, and the application of data-centric techniques and formal methods to the design, analysis and implementation of networked systems. He was awarded the 2006 David J. Sakrison Memorial Prize for the most outstanding dissertation research in the Department of EECS at UC Berkeley, and the 2007 ACM SIGMOD Dissertation Award. He is a recipient of the NSF CAREER award (2009). He has previously served as the program co-chair for the CoNEXT 2008 Student Workshop and the NetDB 2009 workshop co-located with SOSP.

SIGMOD Tutorial 3: Privacy-aware Data Management in Information Networks

Location: Athens View

Presenters: Michael Hay (Cornell University), Kun Liu (Yahoo! Labs), Gerome Miklau (University of Massachusetts Amherst), Jian Pei (Simon Fraser University), Evimaria Terzi (Boston University)

The proliferation of information networks, as a means of sharing information, has raised privacy concerns for enterprises who manage such networks and for individual users that participate in such networks. For enterprises, the main challenge is to satisfy two competing goals: releasing network data for useful data analysis and also preserving the identities or sensitive relationships of the individuals participating in the network. Individual users, on the other hand, require personalized methods that increase their awareness of the visibility of their private information.

Michael Hay received his doctoral dissertation from the University of Massachusetts Amherst in 2010. His thesis is titled "Enabling Accurate Analysis of Private Network Data," and was supervised by David Jensen and Gerome Miklau. He is now a Computing Innovation Fellow at Cornell University under the mentorship of Johannes Gehrke.

Kun Liu, Ph.D., is a scientist at Yahoo! Labs, working on computational advertising with a particular focus on behavioral targeting and social targeting. Prior to that, he was a postdoctoral researcher at IBM Almaden Research Center, working on privacy-preserving social-network analysis, text analytics and healthcare informatics. He has co-authored three book chapters on privacy-preserving social-network analysis and privacy-preserving data mining. His work won him five IBM Invention Achievement Awards, one IBM Invention Plateau Award, one IBM Bravo Award and several Yahoo Display Advertising team awards.

Gerome Miklau is an Assistant Professor of Computer Science at the University of Massachusetts, Amherst. His research focuses on security and privacy challenges in the context of large-scale data management. In recent work he has developed private mechanisms for the analysis of social networks, network traces, and audit logs. He received an NSF CAREER Award in 2007 and won the 2006 ACM SIGMOD Dissertation Award. He received his Ph.D. in Computer Science from the University of Washington in 2005.

Jian Pei is an Associate Professor at the School of Computing Science at Simon Fraser University, Canada. His research interests are to develop effective and efficient data analysis techniques for novel data intensive applications. He is currently interested in various techniques of data mining, Web search, information retrieval, data warehousing, online analytical processing, and database systems, as well as their applications in social networks, health-informatics, business and bioinformatics. His research has been extensively supported in part by many government funding agencies and industry partners. He has published prolifically and served regularly for the leading academic journals and conferences in his fields. He is an associate editor-in-chief of IEEE Transactions of Knowledge and Data Engineering (TKDE), and an associate editor or editorial board member of ACM Transactions on Knowledge Discovery from Data (TKDD), Data Mining and Knowledge Discovery, Statistical Analysis and Data Mining, Intelligent Data Analysis, and Journal of Computer Science and Technology. He is a senior member of the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE). He is the recipient of several prestigious awards.

Evimaria Terzi is an assistant professor in the dept of computer science of Boston University. Prior to joining Boston University in 2009, she was a researcher staff member at IBM Almaden Research Center. Evimaria's work focuses on algorithmic data mining with emphasis on social network and sequence mining. Evimaria got her PhD in 2007 from the University of Helsinki (Finland) and her MSc from Purdue University, USA.

SIGMOD Tutorial 4: Web Data Management

Location: Athens View

Presenters: Michael J. Cafarella (University of Michigan), Alon Y Halevy (Google, Inc.)

Web Data Management (or WDM) refers to a body of work concerned with leveraging the large collections of structured data that can be extracted from the Web. Over the past few years, several research and commercial efforts have explored these collections of data with the goal of improving Web search and developing mechanisms for surfacing different kinds of search answers. This work has leveraged (1) collections of structured data such as HTML tables, lists and forms, (2) recent ontologies and knowledge bases created by crowd-sourcing, such as Wikipedia and its derivatives, DBPedia, YAGO and Freebase, and (3) the collection of text documents from the Web, from which facts could be extracted in a domain-independent fashion. The promise of this line of work is based on the observation that new kinds of results can be obtained by leveraging a huge collection of independently created fragments of data, and typically in ways that are wholly unrelated to the authors' original intent. For example, we might use many database schemas to compute a schema thesaurus. Or we might examine many spreadsheets of scientific data that reveal the aggregate practice of an entire scientific field. As such, WDM is tightly linked to Web-enabled collaboration, even (or especially) if the collaborators are unwitting ones. We will cover the key techniques, principles and insights obtained so far in the area of Web Data Management.

Michael Cafarella is an assistant professor of Computer Science and Engineering at the University of Michigan. His research interests include databases, information extraction, and data mining. He is particularly interested in applying data mining techniques to Web data and scientific applications.

Alon Halevy heads the Structured Data Management Research group at Google. Prior to that, he was a professor of Computer Science at the University of Washington in Seattle, where he founded the database group. In 1999, Dr. Halevy co–founded Nimble Technology, one of the first companies in the Enterprise Information Integration space, and in 2004, Dr. Halevy founded Transformic Inc., a company that created search engines for the deep web, and was acquired by Google. Dr Halevy is a Fellow of the Association for Computing Machinery, received the the Presidential Early Career Award for Scientists and Engineers (PECASE) in 2000, and was a Sloan Fellow (1999–2000). He received his Ph.D in Computer Science from Stanford University in 1993.

SIGMOD Tutorial 5: Data Management Over Flash Memory

Presenters: Ioannis Koltsidas (IBM Research), Stratis D Viglas (University of Edinburgh)

Flash SSDs are quickly becoming mainstream and emerge as alternatives to magnetic disks. It is therefore imperative to incorporate them seamlessly into the enterprise. We present the salient results of research in the area, touching all aspects of the data management stack: from the fundamentals of flash technology, through storage for database systems and the manipulation of SSD-resident data, to query processing.

Ioannis Koltsidas Dr. Koltsidas is a Research Staff Member in the Storage Technologies Department at the Zurich Research Lab. He received a B.S. degree in Electrical and Computer Engineering from the National Technical University of Athens, Greece and a Ph.D. degree in Computer Science from the University of Edinburgh, UK. He subsequently joined IBM at the Zurich Research Lab, where he has worked on flash-based storage, tape storage and distributed file systems.

Stratis D. Viglas is a Reader at the School of Informatics of the University of Edinburgh. He received a PhD in Computer Science from the University of Wisconsin---Madison in 2003, and BSc and MSc degrees in Informatics from the University of Athens, Greece, in 1996 and 1999.

SIGMOD Tutorial 6: Large-Scale Copy Detection

Xin Luna Dong (AT&T Labs - Research), Divesh Srivastava (AT&T Labs - Research)

The Web has enabled the availability of a vast amount of useful information in recent years. However, the web technologies that have enabled sources to share their information have also made it easy for sources to copy from each other and often publish without proper attribution. Understanding the copying relationships between sources has many benefits, including helping data providers protect their own rights, improving various aspects of data integration, and facilitating in-depth analysis of information flow. The importance of copy detection has led to a substantial amount of research in many disciplines of Computer Science, based on the type of information considered, such as text, images, videos, software code, and structured data. This tutorial explores the similarities and differences between the techniques proposed for copy detection across the different types of information. We also examine the computational challenges associated with large-scale copy detection, indicating how they could be detected efficiently, and identify a range of open problems for the community.

Dr. Xin Luna Dong is a researcher at AT&T Labs-Research. She received a Ph.D. in Computer Science and Engineering from University of Washington in 2007, received a Master's Degree in Computer Science from Peking University in China in 2001, and received a Bachelor's Degree in Computer Science from Nankai University in China in 1998. Her research interests include databases, information retrieval and machine learning, with an emphasis on data integration, data cleaning, personal information management, and web search. She has led the Solomon project, whose goal is to detect copying between structured sources and to leverage the results in various aspects of data integration, and the Semex personal information management system, which won the Best Demo award (one of top-3) in Sigmod'05. She co-chaired WebDB'10 and has served in the program committee of Sigmod'11, VLDB'11, PVLDB'10, WWW'10, ICDE'10, VLDB'09, etc.

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. His research interests and publications span a variety of topics in data management.

PODS Tutorial 1: Theory of Data Stream Computing: Where to Go

Location: Horizon

Presenter: S. Muthukrishnan (Rutgers University)

Computing power has been growing steadily, just as communication rate and memory size. Simultaneously our ability to create data has been growing phenomenally and therefore the need to analyze it. We now have examples of massive data streams that are (i) created in far higher rate than we can capture and store in memory economically; (ii) gathered in far more quantity than can be transported to central databases without overwhelming the communication infrastructure; and (iii) arrives far faster than we can compute with them in a sophisticated way.

This phenomenon has challenged how we store, communicate and compute with data. Theories developed over past 50 years have relied on full capture, storage and communication of data. Instead, what we need for managing modern massive data streams are new methods built around working with less. The past 10 years have seen new theories emerge in computing (data stream algorithms), communication (compressed sensing), databases (data stream management systems) and other areas to address the challenges of massive data streams. Still, lot remains open and new applications of massive data streams have emerged recently. We present an overview of these challenges.

S. (Muthu) Muthukrishnan is a Professor in Rutgers Univ. with research interest in databases and algorithms, recently on data stream management and in algorithms for Internet ad systems. This tutorial is based on more than a decade of work on data stream management.

PODS Tutorial 2: Querying Semantic Web Data with SPARQL: State of the Art and Research Perspectives

Location: Horizon

Presenters: Marcelo Arenas (PUC Chile), Jorge Pérez (Universidad de Chile)

The Semantic Web is the initiative of the W3C to make information on the Web readable not only by humans but also by machines. RDF is the data model for Semantic Web data, and SPARQL is the standard query language for this data model. In the last ten years, we have witnessed a constant growth in the amount of RDF data available on the Web, which have motivated the theoretical study of some fundamental aspects of SPARQL and the development of efficient mechanisms for implementing this query language. Some of the distinctive features of RDF have made the study and implementation of SPARQL challenging. First, as opposed to usual database applications, the semantics of RDF is open world, making RDF databases inherently incomplete. Thus, one usually obtains partial answers when querying RDF with SPARQL, and the possibility of adding optional information if present is a crucial feature of SPARQL. Second, RDF databases have a graph structure and are interlinked, thus making graph navigational capabilities a necessary component of SPARQL. Last, but not least, SPARQL has to work at Web-scale! RDF and SPARQL have attracted interest from the database community. However, we think that this community has much more to say about these technologies, and, in particular, about the fundamental database problems that need to be solved in order to provide solid foundations for the development of these technologies. In this tutorial, we will survey some of the main results about the theory of RDF and SPARQL, putting emphasis on some research opportunities for the database community.

Marcelo Arenas is an Associate Professor at the Department of Computer Science at the Pontificia Universidad Catolica de Chile. He received his PhD from the University of Toronto in 2005. His research interests are in different aspects of database theory, such as expressive power of query languages, database semantics, inconsistency handling, database design, XML databases, data exchange, metadata management and database aspects of the Semantic Web. He has received an IBM Ph.D. Fellowship (2004), four best paper awards (PODS 2003, PODS 2005, ISWC 2006 and ICDT 2010) and an ACM-SIGMOD Dissertation Award Honorable Mention in 2006 for his Ph.D dissertation ``Design Principles for XML Data''. He has served on multiple program committees, and since 2009 he has been participating as an invited expert in the World Wide Web Consortium.

Jorge Pérez received a B.Sc. degree in Computer Engineering (2003), a M.Sc. degree in Computer Science (2004), and a Ph.D. degree in Computer Science (2011) from the P. Universidad Católica de Chile. He is currently affiliated with the Computer Science Department at Universidad de Chile. His research interests include data exchange and integration, schema mapping management, and the application of database tools and techniques to the Semantic Web. Pérez has received a Microsoft Research Ph.D. Fellowship (2009-2010), and three best paper awards (ISWC 2006, ESWC 2007, PODS 2011) for his work on database aspects of the Semantic Web, and on extensions of data exchange.

PANELS

Panel 1: Data Management Issues in Health and Medical Informatics

Moderator: Tamer Özsu (University of Waterloo)

Panelists: Allen Brown (Microsoft), Helen Chen (Agfa Healthcare and University of Waterloo), Asuman Dogac (SRDC Ltd), Raymond Ng (UBC), Joel Saltz (Emory University), Wang Chiew Tan (IBM Research – Almaden and University of California – Santa-Cruz)

Health and medical informatics, broadly defined, is the application of computing technologies in healthcare and medicine. These are data intensive domains involving multiple modalities, and there are special constraints on the use of the data. Therefore, the field is rich with data management problems. There are fragmented efforts within data management community in addressing these issues, but, as yet, there is no consistent research agenda. This panel's objective is to start a wide ranging discussion within the community on data management issues in health and medical informatics. We will highlight some research problems and efforts. The panelists represent a balance between academia and industry and between database researchers and domain experts who are involved in research in this area. With participation of others in the audience who are currently involved in this field, we hope to have a lively and engaging discussion of the issues.

Dr. Özsu is a Professor of Computer Science at the David R. Cheriton School of Computer Science of the University of Waterloo. He was the Director of the Cheriton School of Computer Science from January 2007 to June 2010. His research primarily focuses on distributed data management and multimedia data management. He also has an emerging interest in data management issues in health and medical informatics. He holds a PhD from Ohio State University on Computer and Information Science. His B.S. and M.S. are from Middle-East Technical University. He is a Fellow of the ACM and the IEEE.

Since joining Microsoft, Dr. Allen L. Brown, Jr. has been a software architect with a primary focus on distributed systems technologies, including formal semantics, protocols, languages and the development of international standards related to web services. Most recently at Microsoft he has been chief architect of a software platform intended to support investigators and investigations in life sciences. In the years immediately prior to joining Microsoft, he co-founded and developed a venture-backed startup enterprise focused on digital media management. Roughly half of Dr. Brown's thirty-five years of post-Ph.D. involvement in computing was spent at the Xerox Corporation where he served, among other things, as a Research Fellow in its Corporate Research and Technology Division and as CTO of its XSoft Division. While much of the Xerox era was devoted to co-authoring 75+ papers and 24 US patents, Dr. Brown also contributed directly to products, including architecting some of the key facilities of the landmark Xerox "Star" system. He has also held academic appointments as Professor of Computer and Information Science at Syracuse University and James Clark Welling Professor at the George Washington University. Dr. Brown earned undergraduate degrees in Mathematics and Chemical Engineering, and a doctoral degree in Artificial Intelligence, all from the Massachusetts Institute of Technology.

Dr. Helen Chen has been working in the health informatics industry for the past eight years. Prior to her appointment as the Agfa Healthcare Scientist-in-Residence and Assistant Research Professor in the Department of Health Studies and Gerontology with a cross-appointment in the Cheriton School of Computer Science of University of Waterloo, she was a senior researcher at Agfa Healthcare and Agfa's representative at W3C working groups. She is the principle investigator of a number of research projects that apply the emerging web and semantic web technology in the healthcare domain. These projects include "Adaptable Clinical Workflows", "Radiation Dose Registry and Exposure Monitoring", and "Clinical Observation Interoperability". Her research focus is on how to achieve semantic interoperability among complex health information systems during the data exchange and information retrieval. Helen holds a Ph.D in Biomechanics from the University of Waterloo. She received her MSc. And BSc. Degrees from Tsinghua University, Beijing, China.

Dr. Asuman Dogac is the President of SRDC Ltd., which is a spin-off of the Software Research and Development Center, Middle East Technical University. She was a full professor of Department of Computer Engineering at the Middle East Technical University till 2011. She is a graduate of Department of Electrical Engineering, Middle East Technical University. Her research interests include Semantic Web, Internet computing, e-Business, and eHealth. She has more than hundred publications in refereed International journals and conferences which are available from: http://www.srdc.com.tr/ipublications. In 2004, Prof. Dogac received IBM (USA) Faculty award. She is also the recipient of several local awards including 1991 Mustafa Parlar Research Award, 2000 METU Tarik Somer Superior Achievement award, and 2000 Mustafa Parlar Science award. She is on the editorial board of several international journals.

Dr. Raymond Ng is a professor in computer science at UBC. He is internationally renowned for his data mining studies. He has published over 100 journal and conference papers covering a broad range of topics in informatics, data mining and databases. He has won Best Paper awards from the ACM SIGKDD conference on data mining and the ACM SIGMOD conference on database management. For the past decade, Dr Ng has served as a member of the editorial boards of the VLDB Journal and the IEEE Transactions on knowledge and Data Engineering. He was a program co-chair of ACM SIGKDD 2009, a general co-chair of ACM SIGMOD 2008 and a program co-chair of IEEE ICDE 2009. Dr. Ng is a part-time Chief Informatics Officer of the Canadian National Centre of Excellence on the Prevention of Organ Failures (PROOF) whose mission is to discover, develop and commercialize genomics and proteomics biomarker panels for heart, lung and kidney end-stage diseases.. He is a member of the UBC James Hogg iCAPTURE Centre for Cardiovascular and Pulmonary Research, which is part of the Pathology department in the Faculty of Medicine. He is also affiliated with the British Columbia Cancer Research Centre. He has published over 20 papers in medical journals and bioinformatics journals.

Dr. Joel Saltz is Chair of the Department of Biomedical Informatics and Director of the Center for Comprehensive Informatics at Emory University. He holds professorship positions in Biomedical Informatics, Biostatistics and Bioinformatics and Math Computer Science departments at Emory University and Adjunct positions in Computer Science and Computational Science and Engineering at Georgia Tech. His research interests are in development of systems software and data management systems to support data integration, data federation and large scale data management. He has served on the faculty at Yale, University of Maryland, Johns Hopkins and Ohio State Universities. He holds an MD-PhD in Computer Science from Duke University and residency training at Johns Hopkins.

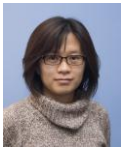

Dr. Wang-Chiew Tan is a Visiting Scientist in the Services Research Department at IBM Research - Almaden and an Associate Professor in the Computer Science Department at University of California, Santa Cruz. She received her B.Sc. (First Class) degree in Computer Science from National University of Singapore, and M.Sc. and Ph.D. degrees in Computer Science from the University of Pennsylvania. Her recent research interests are in database systems, with emphasis on data provenance and information integration. She is the recipient of an NSF CAREER award, a co-recipient of two Best Paper Awards, and several of her publications have been invited to and appeared in special issues for selected papers. For the past year, she has been visiting IBM Almaden Research Center, working on the SPLASH project.

SIGMOD RESEARCH PAPER ABSTRACTS

SIGMOD Research 1: Databases on New Hardware

LazyFTL: A Page-level Flash Translation Layer Optimized for NAND Flash Memory

Dongzhe Ma (Tsinghua University), Jianhua Feng (Tsinghua University), Guoliang Li (Tsinghua University)

Flash is a type of electronically erasable programmable read-only memory (EEPROM), which has many advantages over traditional magnetic disk, such as lower access latency, lower power consumption, lack of noise and shock resistance. Like other EEPROM devices, if a page has been programmed, an erase operation needs to be taken before new data can be written. To make things worse, flash memory can only be erased in blocks which are much larger than the granularity of read and write operations, and therefore flash is also called the write-once bulk-erase device. In addition, after certain amount of erase operations, some blocks of flash memory may become unstable and malfunction. A technology named wear leveling is usually adopted to prolong the life span of flash memory by distributing erase cycles across the entire memory. Flash translation layer (FTL) is a software layer built on raw flash memory which carries out garbage collection and wear leveling strategies and hides the special characteristics of flash memory from upper file systems by emulating a normal block device like magnetic disk. Most existing FTL schemes are optimized for specific access patterns or bring about significant overhead of merge operations under certain circumstances. In this paper, we propose a novel FTL scheme called LazyFTL which acquires low response latency and high scalability, and at the same time, eliminates the overhead of merge operations completely. Our experiment result shows that LazyFTL outperforms all the typical existing FTL schemes today and is very close to the theoretical optimal solution. We also provide a basic design that assists LazyFTL to recover from system failures.

Operation-Aware Buffer Management in Flash-based Systems

Yanfei Lv (Peking University); Bin Cui (Pku); Bingsheng He (Nanyang Technological Universi); Xuexuan Chen (Peking University)

The inherent asymmetry of read and write speeds of flash memory poses great challenges for database system design including buffer management policy. Most of existing flash-based policies adopt the disk-oriented strategies by giving a specific priority to dirty pages, while not fully exploiting the characteristics of flash memory. In this paper, we propose a novel buffer replacement algorithm named FOR, which stands for Flash-based Operation-aware buffer Replacement. The core idea of FOR is based on novel operation-aware page weight determination for buffer replacement. The weight metric not only measures the locality of read/write operations on a page, but also takes the cost difference of read/write operations into account. Moreover, we develop an efficient implementation FOR+ with the time complexity of O(1) for each operation. Experiments with synthetic and benchmark traces demonstrate the efficiency of the proposed strategy, which yields better performance compared with some state-of-the-art flash-based buffer replacement methods.

SkimpyStash: RAM Space Skimpy Key-Value Store on Flash-based Storage

Biplob Debnath (EMC Corporation); Sudipta Sengupta (Microsoft Research); Jin Li (Microsoft Research)

We present SkimpyStash, a RAM space skimpy key-value store on flash-based Storage, designed for high throughput server applications. The distinguishing feature of SkimpyStash is the design goal of extremely low RAM footprint at about 1 (+/- 0.5) byte per key-value pair, which is more aggressive than earlier designs. It uses a combination of linear chaining, bloom filter, and two-choice based balanced allocation policy to tradeoff between RAM space and the low latencies of key-value insert and lookup operations. By being RAM space frugal, SkimpyStash can accommodate larger flash drive capacities for storing and indexing key-value pairs. Our evaluation on commodity server platforms with the real-world datacenter applications traces show that SkimpyStashprovides throughputs from few 10,000s to upwards of 100,000 get-set operations/sec.

Design and Evaluation of Main Memory Hash Join Algorithms for Multi-core CPUs

Spyros Blanas (Univ. of Wisconsin); Yinan Li (Univ. of Wisconsin); Jignesh Patel (Univ. of Wisconsin)

Over the last few years, multi-core chips have dominated the CPU market. However, most database systems have continued to treat multi-core CPUs as less efficient multi-processor systems. This paper investigates how to efficiently execute hash joins in modern multi-core CPUs. We incorporate ideas from the vast literature on parallel join execution in shared-memory and shared-nothing environments and we evaluate them on three radically different modern architectures. Although our investigations are driven by experiments, our analysis reveals that the performance differences are not engineering artifacts, but are tied to intrinsic properties of the memory hierarchy. Surprisingly, we found that the widely used technique of data partitioning often has an adverse impact on overall join performance. Furthermore, frequently examined architectural events, like cache misses, fail to capture the sharing and contention effects that naturally arise in multi-core systems.

SIGMOD Research 2: Query Processing and Optimization

Query Optimization Techniques for Partitioned Tables

Herodotos Herodotou (Duke University); Nedyalko Borisov (Duke University); Shivnath Babu (Duke University)

Table partitioning splits a table into smaller parts that can be accessed, stored, and maintained independent of one another. From their original use to improve query performance, partitioning strategies have evolved into powerful mechanisms to improve database system manageability. Partitioning simplifies administrative tasks like data loading, removal, backup, statistics maintenance, and storage provisioning. SQL extensions and MapReduce frameworks now enable applications and user queries to specify how their results should be partitioned for further use. However, query optimization techniques have not kept pace with the rapid advances in usage and user control of table partitioning. We address this gap by developing new techniques to generate efficient plans for SQL queries involving multiway joins over partitioned tables. Our techniques are designed for easy incorporation into bottom-up query optimizers that are in wide use today. We have prototyped these techniques in the PostgreSQL optimizer. An extensive evaluation shows that our partition-aware optimization techniques, with low optimization overhead, generate plans that can be an order of magnitude better than plans produced by current optimizers.

CrowdDB: Answering Queries with Crowdsourcing

Michael Franklin (UC Berkeley); Donald Kossmann (ETH Zurich); Tim Kraska (UC Berkeley); Sukriti Ramesh (ETH Zurich); Reynold Xin (UC Berkeley)

Some queries cannot be answered by machines only. Processing such queries requires human input for providing information that is missing from the database, for performing computationally-difficult functions and for matching, ranking, or aggregating results based on fuzzy criteria. This paper describes the design and implementation of CrowdDB, a database system that uses human input via crowdsourcing to process queries that neither database systems nor search engines can adequately answer. CrowdDB uses SQL both as a language for posing complex queries and as a way to model data. While CrowdDB leverages many concepts from traditional database systems, there are also important differences. From a conceptual perspective, a major change is that the traditional closed-world assumption for query processing does not hold for human input, which is effectively unbounded. From an implementation perspective, CrowdDB uses an operator-based query engine but this engine is extended with special operators that generate User Interfaces in order to solicit, integrate and cleanse human input. Furthermore, performance and cost depend on a number of different factors including worker affinity, training, fatigue, motivation and location. Real-life experiments conducted using Amazon Mechanical Turk show that CrowdDB is indeed able to process queries that cannot be processed with traditional systems. Furthermore, these experiments assess response time, cost (in $), and result quality depending on a number of different parameters such as the financial reward promised to humans and the way jobs are posted. In this paper, we describe the design of CrowdDB, report on an initial set of performance experiments and outline important avenues for future work in the development of crowdsourced query processing systems.

Skyline Query Processing over Joins

Akrivi Vlachou (NTNU); Christos Doulkeridis (NTNU); Neoklis Polyzotis (UC - Santa Cruz)

This paper addresses the problem of efficiently computing the skyline set of a relational join. Existing techniques either require to access all tuples of the input relations or demand specialized multi-dimensional access methods to generate the skyline join result. To avoid these inefficiencies, we introduce the novel SFSJ algorithm that fuses the identification of skyline tuples with the computation of the join. SFSJ is able to compute the correct skyline set by accessing only a subset of the input tuples, i.e., it has the property of early termination. SFSJ employs standard access methods for reading the input tuples and is readily implementable in an existing database system. Moreover, it can be used in pipelined execution plans, as it generates the skyline tuples progressively. Additionally, we formally analyze the performance of SFSJ and propose a novel strategy for accessing the input tuples that is proven to be optimal for SFSJ. Finally, we present an extensive experimental study that validates the effectiveness of SFSJ and demonstrates its advantages over existing techniques.

Efficient Parallel Skyline Processing using Hyperplane Projections

Henning Koehler (University of Queensland); Jing Yang; Xiaofang Zhou (University of Queensland)

The skyline of a set of multi-dimensional points (tuples) consists of those points for which no clearly better point exists in the given set, using component-wise comparison on domains of interest. As skyline queries can be computationally expensive, it is natural to consider parallelized approaches which make good use of multiple processors. We approach this problem by using hyperplane projections to obtain useful partitions of the data set for parallel processing. These partitions not only ensure small local skyline sets, but enable efficient merging of results as well. Our experiments show that our method consistently outperforms similar approaches for parallel skyline computation, regardless of data distribution, and provides insights on the impacts of different optimization strategies.

SIGMOD Research 3: Schema Mapping and Data Integration

Automatic Discovery of Attributes in Relational Databases

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

In this work we try to cluster relational columns into attributes, i.e., to identify strong relationships between columns based on the common properties and characteristics of the values they contain. For example, identifying whether a certain set of columns refers to telephone numbers versus social security numbers, or names of customers versus names of employees. Traditional relational database schema languages use very limited primitive data types and simple foreign key constraints to express relationships between columns. Object oriented schema languages allow the definition of custom data types, but still, certain relationships between columns might be unknown at design time or they might appear only in a particular database instance. Nevertheless, these relationships are an invaluable tool for schema matching, and generally for better understanding and working with the data. Here, we introduce data oriented solutions (we do not consider solutions that assume the existence of any external knowledge), that use statistical measures to identify strong relationships between the values of a set of columns. Interpreting the database as a graph where nodes correspond to database columns and edges correspond to column relationships, we decompose the graph into connected components and cluster sets of columns into attributes. To test the quality of our solutions, we also provide a comprehensive experimental evaluation using real and synthetic datasets.

Scalable Query Rewriting: A Graph-Based Approach

George Konstantinidis (ISI/USC); Jose Luis Ambite (ISI/USC)

In this paper we consider the problem of answering queries using views, which is important for data integration, query optimization, and data warehouses. We consider its simplest form, conjunctive queries and views, which already is NP-complete. Our context is data integration, so we search for maximally-contained rewritings. By looking at the problem from a graph perspective we are able to gain a better insight and develop an algorithm which compactly represents common patterns in the source descriptions, and (optionally) pushes some computation offline. This together with other optimizations result in an experimental performance about two orders of magnitude faster than current state-of-the-art algorithms, rewriting queries using over 10000 views within seconds.

Leveraging Query Logs for Schema Mapping Generation in U-MAP

Hazem Elmeleegy (Purdue University); Ahmed Elmagarmid (Qatar Computing Research Institute); Jaewoo Lee (Purdue University)

In this paper, we introduce U-MAP, a new system for schema mapping generation. U-MAP builds upon and extends existing schema mapping techniques. However, it mitigates some key problems in this area, which have not been previously addressed. The key tenet of U-MAP is to exploit the usage information extracted from the query logs associated with the schemas being mapped. We describe our experience in applying our proposed system to realistic datasets from the retail and life sciences domains. Our results demonstrate the effectiveness and efficiency of U-MAP compared to the traditional approaches.

Designing and Refining Schema Mappings via Data Examples

Bogdan Alexe (UC Santa Cruz); Balder ten Cate (UC Santa Cruz); Phokion Kolaitis (UC Santa Cruz); Wang-Chiew Tan (IBM Research - Almaden and UC Santa Cruz)

A schema mapping is a specification of the relationship between a source schema and a target schema. Schema mappings are fundamental building blocks in data integration and data exchange and, as such, obtaining the right schema mapping constitutes a major step towards the integration or exchange of data. Up to now, schema mappings have typically been specified manually or have been derived using mapping-design systems that automatically generate a schema mapping from a visual specification of the relationship between two schemas. We present a novel paradigm and develop a system for the interactive design of schema mappings via data examples. Each data example represents a partial specification of the semantics of the desired schema mapping. At the core of our system lies a sound and complete algorithm that, given a finite set of data examples, decides whether or not there exists a GLAV schema mapping (i.e., a schema mapping specified by Global-and-Local-As-View constraints) that "fits" these data examples. If such a fitting GLAV schema mapping exists, then our system constructs the "most general" one. We give a rigorous computational complexity analysis of the underlying decision problem concerning the existence of a fitting GLAV mapping, given a set of data examples. Specifically, we prove that this problem is complete for the second level of the polynomial hierarchy, hence, in a precise sense, harder than NP-complete. This worst-case complexity analysis notwithstanding, we conduct an experimental evaluation of our prototype implementation that demonstrates the feasibility of interactively designing schema mappings using data examples. In particular, our experiments show that our system achieves very good performance in real-life scenarios.

SIGMOD Research 4: Data on the Web

Apples and Oranges: A Comparison of RDF Benchmarks and Real RDF Datasets

Songyun Duan (IBM Research - Thomas J. Watson Research Ctr.), Anastasios Kementsietsidis (IBM T.J. Watson), Kavitha Srinivas (IBM T. J. Watson), Octavian Udrea (IBM T. J. Watson)

The widespread adoption of the Resource Description Framework (RDF) for the representation of both open (web) and enterprise data is the driving force behind the increasing research interest in RDF data management. As RDF data management systems proliferate, so do benchmarks to test the scalability and performance of these systems under various data and query workloads. In this paper, we focus on the existing benchmarks, and we offer a detailed comparison between benchmark-generated data and data found in real, and widely used RDF datasets. The results of our comparison illustrate that existing benchmark data have little in common with real data. Therefore any conclusions drawn from existing benchmark tests might not actually translate to expected behaviors in real settings. In terms of the comparison itself, we show that simple primitive data metrics are inadequate in and by themselves in terms of fleshing out the differences between real and synthetic data. In response to the limitations of these metrics, and as part of our first motivation, we introduce intuitive and novel metrics that can indeed highlight the key differences between distinct datasets. In response to the limitations of existing benchmarks, and as part of our second contribution, we introduce a new benchmark generator with the following novel characteristics: (a) the generator can use any (real or synthetic) dataset and convert it into a benchmark dataset; (b) the generator can generate data that mimic the characteristics of real datasets. On the technical side, to create a benchmark generator with the above characteristics, we formulate the benchmark generation problem as a linear programming problem whose solution provides us with the desired benchmark datasets. To our knowledge, this is the first methodological study of RDF benchmarking, as well as the first attempt on generated RDF benchmarks in a principled way.

Efficient Query Answering in Probabilistic RDF Graphs

Xiang Lian (HKUST), Lei Chen (Hong Kong University of Science and Technology)

In this paper, we tackle the problem of efficiently answering queries on probabilistic RDF data graphs. Specifically, we model RDF data by probabilistic graphs, and an RDF query is equivalent to a search over subgraphs of probabilistic graphs that have high probabilities to match with a given query graph. To efficiently process queries on probabilistic RDF graphs, we propose effective pruning mechanisms, structural and probabilistic pruning. For the structural pruning, we carefully design synopses for vertex/edge labels by considering the label distributions, as well as other structural information, in order to improve the pruning power. For the probabilistic pruning, we derive a cost model to guide the pre-computation of probability upper bounds such that the query cost is expected to be low. We construct an index structure that integrates synopses/statistics for both structural and probabilistic pruning, and give an efficient query answering approach to answer queries on probabilistic RDF graph data. The efficiency of our proposed solutions has been verified through extensive experiments.

Facet Discovery for Structured Web Search: A Query-log Mining Approach

Jeffrey Pound (University of Waterloo), Stelios Paparizos (Microsoft Research), Panayiotis Tsaparas

In recent years there has been a strong trend of incorporating results from structured data sources into keyword-based web search systems such as Google or Amazon. When presenting structured data, facets are a powerful tool for navigating, refining, and grouping the results. For a given structured data source, a fundamental problem in supporting faceted search is finding an ordered selection of attributes and values that will populate the facets. This creates two sets of challenges: first, the screen real-estate is limited while there are tens of different attributes available for a given entity type, hence it is important that the top few facets match the user intent to avoid a poor experience. Second, the huge scale of data available to search engines like Google or Amazon demands an automated unsupervised solution. In this paper, we model the user faceted-search behavior using the intersection of web query logs with existing structured data. Since web queries are formulated as free-text queries, a challenge in our approach is the inherent ambiguity in mapping keywords to the different possible attributes of a given entity type. We present an automated solution that elicits user preferences on attributes and values, employing different disambiguation techniques ranging from simple keyword matching, to more sophisticated probabilistic models. We demonstrate experimentally the scalability of our solution by running it on hundreds of categories of diverse entity types, and a large query log. We also present an evaluation of the quality of our results based on a user-study with Mechanical Turk.

Schema-As-You-Go: On Probabilistic Tagging and Querying of Wide Tables

Meiyu Lu (NUS), Bing Tian Dai (NUS), Anthony Tung, Divyakant Agrawal (UC Santa Barbara)

The emergence of Web 2.0 has resulted in a huge amount of heterogeneous data that are contributed by a large number of users, engendering new challenges for data management and query processing. Given that the data are unified from numerous sources and can be accessed by a large number of users, providing users with a unified mediated schema as data integration is insufficient. This is because, on one hand, a deterministic mediated schema restricts user's freedom to express their queries in their preferred vocabulary; on the other hand, it is not realistic for users to remember the numerous attribute names that arise from integrating numerous data sources. As such, a user-oriented data management and query interface is required. In this paper, we propose an out-of-the-box approach that separates users' actions from database

operations. This separating layer deals with the challenges from a semantic perspective. It interprets the semantics of each data value through tags that are provided by users, and then inserts the value into the database together with these tags. When querying the database, this layer also serves as a platform for retrieving data, inferring the semantics of queried tags from the users before providing the query result. Experiments are conducted to illustrate both the effectiveness and efficiency of our approach.

SIGMOD Research 5: Data Privacy and Security

No Free Lunch in Data Privacy

Daniel Kifer (Penn State), Ashwin Machanavajjhala (Yahoo)

Differential privacy guarantees that an attacker who knows every record but one in a statistical database cannot use query answers to infer much about the remaining record because the distribution of noisy query answers has little dependence on whether or not that record has been deleted from the database. Because of this guarantee, there is a widely held belief that differential privacy offers privacy protection without making any assumptions about how the data are generated. In this paper we show that this belief is incorrect. We prove a stronger result that we call the no free lunch theorem: any privacy definition with minimal utility guarantees must make assumptions about the data-generating mechanism. The proof is trivial. The difficulty lies in explaining how differential privacy makes such assumptions. The assumption behind differential privacy is that protecting against attackers who know all but one record provides meaningful privacy guarantees for all data-generating mechanisms. We show that attackers with greater amounts of knowledge do not always pose the greatest privacy risks; we present variations of differential privacy where the attackers are strictly stronger yet less noise is required when answering queries. We then argue that a meaningful privacy guarantee arises not when we can hide the presence of a record, but when we can hide the participation of an individual (or other entity). We show that hiding a record and hiding participation are completely different concepts when there are correlations or causal relations between records. We therefore conjecture that differential privacy assumes that records are independent (but not identically distributed). We illustrate these ideas in the cases of social network models and tabular data (with a prior release of global statistics). In all of these cases, the notion of participation is different and is not adequately protected by differential privacy.

TrustedDB: A Trusted Hardware based Database with Privacy and Data Confidentiality

Sumeet Bajaj (Stony Brook University), Radu Sion (Stony Brook University)

TrustedDB is an outsourced database prototype that allows clients to execute SQL queries with privacy and under regulatory compliance constraints without having to trust the service provider. TrustedDB achieves this by leveraging server-hosted tamper-proof trusted hardware in critical query processing stages. TrustedDB does not limit the query expressiveness of supported queries. And, despite the cost overhead and performance limitations of trusted hardware, the costs per query are orders of magnitude lower than any (existing or) potential future cryptography-only mechanisms. TrustedDB is built and runs on actual hardware, and its performance and costs are evaluated here.

Differentially Private Data Cubes: Optimizing Noise Sources and Consistency

Bolin Ding (UIUC), Marianne Winslett (UIUC), Jiawei Han (UIUC); Zhenhui Li (UIUC)

Data cubes play an essential role in data analysis and decision support. In a data cube, data from a fact table is aggregated on subsets of the table's dimensions, forming a collection of smaller tables called cuboids. When the fact table includes sensitive data such as salary or diagnosis, publishing (even only some of) cuboids may compromise individuals' privacy. In this paper, we address this problem using differential privacy (DP), which provides provable privacy guarantees for individuals by adding noise to query answers. We choose an initial subset of cuboids to compute directly from the fact table, injecting DP noise as usual; and then compute the remaining cuboids directly from the initial set. Given a fixed privacy guarantee, we show that it is NP-hard to choose the initial set so that the maximal noise over all published cuboids is minimized, or so that the number of cuboids with maximal noise below a given threshold (precise cuboids) is maximized. We provide an efficient procedure with running time polynomial in the number of cuboids to select the initial set of cuboids, such that the maximal noise in all published cuboids will be within a factor (ln|lat|+1)^2 of optimal, where |lat| is the number of cuboids to be published, or the number of precise cuboids will be within a factor (1-1/e) of optimal. We also introduce how to enforce consistency in the published cuboids while even improving their utility (reducing error). Through an empirical evaluation on real and synthetic data, we show that our approach is effective by reporting the amount of errors incurred by different publishing algorithms.

iReduct: Differential Privacy with Reduced Relative Errors

Xiaokui Xiao (Nanyang Technological Univ), Gabriel Bender (Cornell University), Michael Hay (Cornell University), Johannes Gehrke (Cornell University)

Differential privacy is the state-of-the-art paradigm for privacy preserving data publishing. The existing techniques for enforcing differential privacy, however, optimize only the worst-case noise variance in the published data. Specifically, given a set Q of queries, those techniques produce a noisy answer for each query in Q, such that the variance of the noise in each result is bounded by a threshold. In other words, different query results may be injected with equal amount of noise, regardless of whether the results are large or small. As a consequence, for those queries whose exact answers are small, their

noisy results tend to be dominated by noise, which leads to inferior data utility. This paper addresses the problem with iReduct, a general framework for generating differentially private results with reduced relative errors. The basic idea of iReduct is to apply different amount of noise in different query results, such that a smaller (larger) result is more likely to be injected with less (more) noise. The backbone of the framework is a novel extension of Dwork et al.'s method that employs correlated Laplace noise to improve data utility. To demonstrate the application of iReduct, we present an instantiation of iReduct for generating marginals, i.e., the projections of a multi-dimensional dataset on various subsets of its attributes. Extensive experiments with real data demonstrate the effectiveness of our solution.

SIGMOD Research 6: Data Consistency and Parallel DB

Fast Checkpoint Recovery Algorithms for Frequently Consistent Applications

Tuan Cao (Cornell University), Marcos Vaz Salles (Cornell University), Benjamin Sowell (Cornell University), Yao Yue (Cornell University), Johannes Gehrke (Cornell University), Alan Demers (Cornell University), Walker White (Cornell University)

Advances in hardware have enabled many long-running applications to execute entirely in main memory. As a result, these applications have increasingly turned to database techniques to ensure durability in the event of a crash. However, many of these applications, such as massively multiplayer online games and main-memory OLTP systems, must sustain extremely high update rates -- often hundreds of thousands of updates per second or more. Providing durability for these applications without introducing excessive overhead or latency spikes remains a challenge for application developers. In this paper, we take advantage of frequent points of consistency in many of these applications to develop novel checkpoint recovery algorithms that trade additional space in main memory for significantly lower overhead and latency. Compared to previous work, our new algorithms do not require any locking or bulk copies of the application state. Our experimental evaluation shows that one of our new algorithms attains nearly constant latency and reduces overhead by more than an order of magnitude for low to medium update rates. Additionally, in a heavily loaded main-memory transaction processing system, it still reduces overhead by more than a factor of two.

A Latency and Fault-Tolerance Optimizer for Online Parallel Query Plans

Prasang Upadhyaya (University of Washington), YongChul Kwon (University of Washington), Magdalena Balazinska (University of Washington)

We address problem of making parallel, pipelined query plans fault-tolerant. Our goal is to maintain pipelining to enable system to produce results incrementally, while providing intra-query fault-tolerance to avoid entire query restarts when failures occur. Our key hypothesis is that parallel system requires not only efficient fault-tolerance strategy that does not require blocking but should also support use of different strategies at different operators within single query plan. Enabling each operator to use different fault tolerance strategy leads to space of possible fault tolerance plans amenable to cost-based optimization. To test our hypothesis, we first develop framework that enables mixing and matching of fault-tolerance techniques in single parallel, pipelined query plan. Second, we develop FTOpt, cost-based fault-tolerance optimizer that, given failure model, automatically selects best strategy for each operator in query plan in manner that minimizes expected processing time with failures for entire query. We implement our approach in prototype parallel query-processing engine. Our experiments demonstrate that (1) there is no single best fault-tolerance strategy for all query plans, (2) often hybrid strategies that mix-and-match different recovery techniques outperform any uniform strategy, and (3) our optimizer is able to correctly identify winning fault-tolerance configurations.

ArrayStore: A Storage Manager for Complex Parallel Array Processing

Emad Soroush (Univ. of Washington), Magdalena Balazinska (Univ. of Washington), Daniel Wang (SLAC)

We present the design, implementation, and evaluation of ArrayStore, a new storage manager for complex, parallel array processing. ArrayStore builds on prior work in the area of multidimensional data storage, but considers the new problem of supporting a parallel and more varied workload comprising not only range-queries, but also binary operations such as joins and complex user-defined functions. This paper makes two key contributions. First, it examines several existing single-site storage management strategies and array partitioning strategies to identify which combination is best suited for the array-processing workload above. Second, it develops a new and efficient storage-management mechanism that enables parallel processing of operations that must access data from adjacent partitions. We evaluate ArrayStore on over 80GB of real data from two scientific domains and real operators used in these domains. We show that ArrayStore outperforms previously proposed storage management strategies in the context of its diverse target workload.

Warding off the Dangers of Data Corruption with Amulet

Nedyalko Borisov (Duke University), Shivnath Babu (Duke University), Nagapramod Mandagere (IBM Almaden Research Center), Sandeep Uttamchandani (IBM Almaden Research Center)

Data corruption is an unfortunate byproduct of the complexity of modern systems. Hardware errors, software bugs, and mistakes by human administrators can corrupt important sources of data. Today, administrators define ad-hoc scripts to verify data consis- tency by running data integrity tests at the application, database, file-system, and storage levels. This

manual approach is tedious and provides no understanding of the potential system unavailabil- ity and data loss in case a corruption occurs. We introduce the Amulet system that addresses the problem of proactively verifying data consistency. To our knowledge, Amulet is the first tool that: (i) automates the data verification process; (ii) considers and ensures that the administrator objectives are met; and (iii) provides timely notifications when problems are detected, allowing the correction process to be proactive Amulet takes the administrator's objectives and searches a multidimensional space to find a testing plan that will satisfy all the objectives specified. This plan is executed con- tinuously, and timely notifications are generated when problems are detected. Amulet was implemented and validated for a database software stack deployed on the Amazon cloud. We demonstrate the effectiveness of Amulet in comparison to a brute-force optimizer.

SIGMOD Research 7: Service Oriented Computing, Data Management in the Cloud

Schedule Optimization for Data Processing Flows on the Cloud

Herald Kllapi (University Of Athens), Eva Sitaridi (University of Athens), Manolis Tsangaris (University of Athens), Yannis Ioannidis (University of Athens)

Scheduling data processing workflows (dataflows) on the cloud is a very complex and challenging task. It is essentially an optimization problem, very similar to query optimization, that is characteristically different from traditional problems in two aspects. First, its space of alternative schedules is very rich, due to various optimization opportunities that cloud computing offers. Second, its optimization criterion is at least two-dimensional, with the monetary cost of using the cloud being at least as important as the query completion time. In this paper, we study the scheduling of dataflows that involve arbitrary data processing operators in the context of three different problems: 1) minimize completion time given a fixed budget, 2) minimize monetary cost given a time limit, and 3) find trade-offs between completion time and monetary cost without any a-priori constraints. We formulate the problems and present an approximate optimization framework to address them that makes use of resource elasticity in the cloud. To investigate the effectiveness of our approach, we incorporate the devised framework into a prototype system for dataflow evaluation and instantiate it with several greedy, probabilistic, and exhaustive search algorithms. Finally, we present the results of several experiments that we have conducted with the prototype elastic optimizer on numerous scientific and synthetic dataflows and we identify the advantages and disadvantages of the various search algorithms. The overall results are quite promising and indicate the effectiveness of our approach.

Zephyr: Live Migration in Shared Nothing Databases for Elastic Cloud Platforms

Aaron Elmore (UC Santa Barbara), Sudipto Das (UC Santa Barbara), Divyakant Agrawal (UC Santa Barbara), Amr El Abbadi (UC Santa Barbara)

Efficient and scalable multitenant application and data infrastructures are crucial to the success of large cloud platforms hosting hundreds of thousands of applications. A majority of these applications are characterized by a small data footprint and unpredictable load patterns. When such a platform is built on an elastic pay-peruse infrastructure, an added challenge is to minimize the system's operating cost while guaranteeing the tenants' service level agreements (SLA). Elastic load balancing is therefore an important feature to enable scale up during high load while scaling down when the load is low. Live migration, a technique to migrate live tenants with minimal service interruption and no downtime, is critical to allow lightweight elastic scaling. We focus on the problem of live migration in the database layer. We propose Zephyr, a technique to efficiently migrate a tenant's live database in a shared nothing transactional database architecture. Zephyr uses phases of on-demand pull and asynchronous push of data, requires minimal synchronization, results in minimal service interruption with few or no aborted transactions, minimizes the amount of data transferred, and continues to provide ACID guarantees during migration. We provide a systematic analysis of the guarantees provided and prove correctness of the design. A prototype implementation using an open source relational database engine and experiments using transaction processing benchmarks demonstrates its efficiency. Migrating a live database using Zephyr results in only a few tens of failed operations, 10-20% change in average transaction latency, minimal messaging, and no overhead during normal operation.

Workload-Aware Database Monitoring and Consolidation

Carlo Curino (MIT), Evan Jones (MIT), Sam Madden (MIT), Hari Balakrishnan (MIT)

In most enterprises, databases are deployed on dedicated database servers. Often, these servers are underutilized much of the time. For example, in traces from almost 200 production servers from different organizations, we see an average CPU utilization of less than 4%. This unused capacity can be harnessed to consolidate multiple databases on a fewer number of machines, reducing hardware and operational costs. Virtual machine (VM) technology is one popular way to approach this problem. However, as we demonstrate in this paper, VMs fail to adequately support database consolidation, because databases place a unique and challenging set of demands on hardware resources, which are not well-suited to the assumptions made by VM-based consolidation. Our system for database consolidation, named Kairos, uses novel techniques to measure the hardware requirements of database workloads, as well as models to predict the combined resource utilization of those workloads. We formalize the consolidation problem as a non-linear optimization program, aiming to minimize the number of servers and balancing load, while achieving near-zero degradation of transactional

performance of the consolidated workload. We compare Kairos against various VM-based approaches, showing up to a factor 12× performance improvement on a TPC-C-like benchmark. We also tested the effectiveness of our approach on real-world data collected from production servers at Wikia.com, Wikipedia, Second Life, and our institution, showing absolute consolidation factors ranging between 6:1 and 17:1.

Predicting Cost Amortization for Query Services

Verena Kantere (Cyprus University of Technology), Debabrata Dash (ArcSight), Georgios Gratsias (ELCA Informatique SA), Anastasia Ailamaki (EPFL)

Emerging providers of online services offer access to data collections. Query service providers need to build data structures, such as materialized views and indexes, in order to offer better performance for user query execution. The cost of such structures is charged to the user as part of the overall query service cost. In order to ensure the economic viability of the provider, the building and maintenance cost of new structures has to be amortized to a set of prospective query services that will use them. This work proposes a novel stochastic model that predicts the extent of cost amortization in time and number of services. The model is completed with a novel method that regresses query traffic statistics and provides input to the prediction model. In order to demonstrate the effectiveness of the prediction model, we study its application on an extension of an existing economy model for the management of a cloud DBMS. A thorough experimental study shows that the prediction model ensures the economic viability of the cloud DBMS while enabling the offer of fast and cheap query services.

Performance Prediction for Concurrent Database Workloads

Jennie Duggan (Brown University), Ugur Cetintemel (Brown University), Olga Papaemmanouil (Brandeis University), Eli Upfal (Brown University)

Current trends in data management systems, such as cloud and multi-tenant databases, are leading to data processing environments that concurrently execute highly heterogeneous query workloads while at the same time need to satisfy diverse performance expectations. In these newly-arising settings avoiding potential Quality-of-Service (QoS) violations strongly relies on performance predictability, i.e., being able to estimate the impact of concurrent query execution on the performance of each individual query in an ever-changing workload mix. In this work we present an approach based on resource contention modeling to estimate the impact of concurrency on analytical query performance. Specifically, our solution relies on the analysis of query behavior in isolation, pairwise interactions and sampling techniques to predict the resource contention under various query mixes and concurrency levels. Resource contention predictions are then used to obtain an estimation of the end-to-end query latency. To improve the accuracy of the predictions, our framework uses the resource contention model to create a timeline of the query interactions, i.e., a fine-grained estimation of the time segments during which discrete mixes will be executed concurrently. An experimental evaluation on TPC-H queries demonstrates that our models can provide predictions within approximately 20% of the actual values in the average case.

SIGMOD Research 8: Spatial and Temporal Data Management

Reverse Spatial and Textual k Nearest Neighbor Search

Jiaheng Lu (Renmin University of China), Ying Lu (Renmin University of China), Gao Cong (Nanyang Technological University)

Geographic objects associated with descriptive texts are becoming prevalent. This gives prominence to spatial keyword queries that take into account both the locations and textual descriptions of content. Specifically, the relevance of an object to a query is measured by spatial-textual similarity that is based on both spatial proximity and textual similarity. In this paper, we define Reverse Spatial Textual k Nearest Neighbor (RSTkNN) query, i.e., finding objects that take the query object as one of their k most spatial-textual similar objects. Existing work on reverse kNN queries focus solely on spatial locations but ignore textual relevance. To answer RSTkNN queries efficiently, we propose a hybrid index tree called IUR-Tree (Intersection-Union R-Tree) that effectively combines location proximity with textual document similarity. Based on the IUR-Tree, we design a branch-and-bound search algorithm. To further accelerate the query processing, we propose an enhanced variant of the IUR-Tree called clustered IUR-Tree and corresponding optimization algorithms. Empirical studies show that the proposed algorithms offer scalability and are capable of excellent performance.

Location-Aware Type Ahead Search on Spatial Databases : Semantics and Efficiency

Senjuti Basu Roy (UT Arlington), Kaushik Chakrabarti (Microsoft Research)

Spatial databases like yellow page data are often searched using keywords. For example, consider a user searching for a Starbucks or a Chinese restaurant from a mobile phone. Such systems are ``location-aware", i.e., the spatial objects are ranked not only by relevance to the keyword query but also by their proximity to user's location and their static scores (e.g., popularity, ratings). Typing the entire query is cumbersome and prone to errors, especially from mobile phones. We address this problem by introducing type-ahead search functionality for spatial databases. Like keyword search on spatial data, a type-ahead search needs to be location-aware, i.e., with every letter being typed, it needs to return spatial objects whose prefix (either on names or descriptions) contains the query string typed so far and which ranks highest in terms of proximity and other static scores. Existing solutions for type-ahead search cannot be directly used as they are not location-

aware. We show that a straight-forward combination of existing techniques for performing type-ahead search with those for performing proximity search perform poorly. We propose a formal model for query processing cost and develop novel techniques that optimize that cost. Our empirical evaluation on real and synthetic datasets demonstrate the effectiveness of our techniques. To the best of our knowledge, this is the first work on location-aware type-ahead search.

Collective Spatial Keyword Querying

Xin Cao (Singapore NTU), Gao Cong (Nanyang Technological University), Christian Jensen (Aarhus University), Beng Chin Ooi (National University of Singapore)

With the proliferation of geo-positioning and geo-tagging, spatial web objects that posses both a geographical location and a textual description are gaining in prevalence, and spatial keyword queries that exploit both location and textual description are gaining in prominence. However, the queries studied so far generally focus on finding individual objects that each satisfy a query rather than finding groups of objects where the objects in a group collectively satisfy a query. We define the problem of retrieving a group of spatial web objects such that the group's keywords cover the query's keywords and such that objects are nearest to the query location and have the lowest inter-object distances. Specifically, we study two variants of this problem, both of which are NP-complete. We devise exact solutions as well as approximate solutions with provable approximation bounds to the problems. We present empirical studies that offer insight into the efficiency and accuracy of the solutions.

Finding Semantics in Time Series

Peng Wang (Fudan University), Haixun Wang (Microsoft), Wei Wang (Fudan University)

In order to understand a complex system, we start by analyzing its output or its log data. For example, we track a system's resource consumption (CPU, memory, message queues of different types, etc) to help avert system failures; we examine economic indicators to assess the severity of a recession; we monitor a patient's heart rate or EEG for disease diagnosis. Time series data is involved in many such applications. Much work has been devoted to pattern discovery from time series data, but not much has attempted to use the time series to unveil a system's internal dynamics. In this paper, we go beyond learning patterns from time series data. We focus on obtaining a better understanding of its data generating mechanism, and we regard patterns and their temporal relations as organic components of the hidden mechanism. Specifically, we ropose to model time series data using a novel pattern-based hidden Markov model pHMM), which aims at revealing a global picture of the system that generates the time series data. We propose an iterative approach to refine pHMMs learned from the data. In each iteration, we use the current pHMM to guide time series segmentation and clustering, which enables us to learn a more accurate pHMM. Furthermore, we propose three pruning strategies to speed up the refinement process. Empirical results on real datasets demonstrate the feasibility and effectiveness of the proposed approach.

Querying Contract Databases Based on Temporal Behavior

Elio Damaggio (UCSD); Alin Deutsch (UCSD); Dayou Zhou (UCSD)

Considering a broad definition for service contracts (beyond web services and software, e.g. airline tickets and insurance policies), we tackle the challenges of building a high performance broker in which contracts are both specified and queried through their temporal behavior. The temporal dimension, in conjunction with traditional relational attributes, enables our system to better address difficulties arising from the great deal of information regarding the temporal interaction of the various events cited in contracts (e.g. "No refunds are allowed *after* a reschedule of the flight, which can be requested only *before* any flight leg has been used"). On the other hand, querying large repositories of temporal specifications poses an interesting indexing challenge. In this paper, we introduce two distinct and complementary indexing techniques that enable our system to scale the evaluation of a novel and theoretically sound notion of 'permission' of a temporal query by a service contract. Our notion of permission is inspired by previous work on model checking but, given the specific characteristic of our problem, does not reduce to it. We evaluate experimentally our implementation, showing that it scales well with both the number and the complexity of the contracts.

SIGMOD Research 9: Shortest Paths and Sequence Data

Neighborhood-Privacy Protected Shortest Distance Computing in Cloud

Jun Gao (Peking University), Jeffrey Xu Yu (The Chinese University of Hong Kong), Ruoming Jin (Kent State University), Jiashuai Zhou, Tengjiao Wang, Dongqing Yang

With the advent of the cloud computing, it becomes desirable to utilize the cloud computing to efficiently process complex operations in large graphs without compromising their sensitive information. This paper studies shortest distance computing in the cloud, which aims at the following goals, i) preventing outsourced graphs from neighborhood attack, ii) preserving the shortest distance in outsourced graphs, iii) minimizing the overhead on the client side. The basic idea of this paper is to transform an original graph G into a link graph Gl kept locally and a set of outsourced graphs Go. Each outsourced graph should meet the requirement of a new security model called 1-neighborhood-d-radius. In addition, the shortest distance query can be equivalently answered using Gl and Go. We show that the optimal graph transformation which minimizes the overhead on the client side with the satisfaction of both security and utility requirements is NP-Hard. We devise a greedy method to produce Go and Gl, which can exactly answer the shortest distance queries. We also develop an efficient

transformation method to support approximate shortest distance answering under a given additive error bound. The final results illustrate the effectiveness and efficiency of our method.

On k-skip Shortest Paths

Yufei Tao (Chinese Univ. of Hong Kong), Cheng Sheng (The Chinese University of Hong Kong), JIan Pei (Simon Fraser Univ.)

Given two vertices s, t in a graph, let P be the shortest path (SP) from s to t, and P* a subset of the vertices in P. P* is a k-skip shortest path from s to t, if it includes at least a vertex out of every k consecutive vertices in P. In general, P* succinctly describes P by sampling the vertices in P with a rate of at least 1/k. This makes P* a natural substitute in scenarios where reporting every single vertex of P is unnecessary or even undesired. This paper studies k-skip SP computation in the context of spatial network databases (SNDB). Our technique has two properties crucial for real-time query processing in SNDB. First, our solution is able to answer k-skip queries significantly faster than finding the original SPs in their entirety. Second, the previous objective is achieved with a structure that occupies less space than storing the underlying road network. The proposed algorithms are the outcome of a careful theoretical analysis that reveals valuable insight into the characteristics of the k-skip SP problem. Their efficiency has been confirmed by extensive experiments with real data.

Finding Shortest Path on Land Surface

Lian Liu (HKUST); Raymond Chi-Wing Wong (The Hong Kong University of Sc)

Finding shortest paths is a fundamental operator in spatial databases. Recently, terrain datasets have attracted a lot of attention from both industry and academia. There are some interesting issues to be studied in terrain datasets which cannot be found in a traditional two-dimensional space. In this paper, we study one of the issues called a slope constraint which exists in terrain datasets. In this paper, we propose a problem of finding shortest paths with the slope constraint. Then, we show that this new problem is more general than the traditional problem of finding shortest paths without considering the slope constraint. Since finding shortest paths with the slope constraint is costly, we propose a new framework called surface simplification so that we can compute shortest paths with the slope constraint efficiently. We conducted experiments to show that the surface simplification is very efficient and effective not only for the new problem with the slope constraint but also the traditional problem without the slope constraint.

WHAM: A High-throughput Sequence Alignment Method

Yinan Li (Univ. of Wisconsin), Allison Terrell (Univ. of Wisconsin), Jignesh Patel (Univ. of Wisconsin)

Over the last decade the cost of producing genomic sequences has dropped dramatically due to the current so called "next-gen" sequencing methods. However, these next-gen sequencing methods are critically dependent on fast and sophisticated data processing methods for aligning a set of query sequences to a reference target genome using rich string matching models. The focus of this work is on the design, development and evaluation of a database index-based method for this crucial "short read alignment" problem. Our method, called WHAM, employs novel hash-based indexing methods and bitwise operations for sequence alignments. It allows richer match models than existing methods and it is significantly faster than the existing state-of-the-art method. In addition, its relative speedup over the existing method is poised to increase in the future in which read sequence lengths will increase.

A New Approach for Processing Ranked Subsequence Matching Based on Ranked Union

Wook-Shin Han (Kyungpook National University), Jinsoo Lee (Kyungpook National University), Yang-Sae Moon, Seung-won Hwang, Hwanjo Yu (POSTECH)

Ranked subsequence matching finds top-k subsequences most sim-ilar to a given query sequence from data sequences. Recently, Han et al. [12] proposed a solution (referred to here as HLMJ) to this problem by using the concept of the minimum distance matching window pair (MDMWP) and a global priority queue. By using the concept of MDMWP, HLMJ can prune many unnecessary accesses to data subsequences using a lower bound distance. However, we notice that HLMJ may incur serious performance overhead for important types of queries. In this paper, we propose a novel systematic framework to this problem by viewing ranked subsequence matching as ranked union. Specifically, we propose a notion of the matching subsequence equivalence class (MSEQ) and a novel lower bound called the MSEQ-distance. To completely eliminate the performance problem of HLMJ, we also propose a cost-aware density-based scheduling technique, where we consider both density and cost of the priority queue. Extensive experimental results with many real datasets show that the proposed algorithm outperforms HLMJ and the adapted PSM [23], a state-of-the-art index-based merge algorithm supporting non-monotonic distance functions, by up to 980.9 and 5352.1 times, respectively.

SIGMOD Research 10: Data Provenance, Workflow and Cleaning

Interaction between Record Matching and Data Repairing

Wenfei Fan (University of Edinburgh), Jianzhong Li (Harbin Institute of Technology), Shuai Ma (Beihang University), Nan Tang (University of Edinburgh), Wenyuan Yu (University of Edinburgh)

Central to a data cleaning system are record matching and data repairing. Matching aims to identify tuples that refer to the same real-world object, and repairing is to make a database consistent by fixing errors in the data by using constraints. These are treated as separate processes in current data cleaning systems, based on heuristic solutions. This paper studies a new problem, namely, the interaction between record matching and data repairing. We show that repairing can effectively help us identify matches, and vice versa. To capture the interaction, we propose a uniform framework that seamlessly unifies repairing and matching operations, to clean a database based on integrity constraints, matching rules and master data. We give a full treatment of fundamental problems associated with data cleaning via matching and repairing, including the static analyses of constraints and rules taken together, and the complexity, termination and determinism analyses of data cleaning. We show that these problems are hard, ranging from NP- or coNP-complete, to PSPACE-complete. Nevertheless, we propose efficient algorithms to clean data via both matching and repairing. The algorithms find deterministic fixes and reliable fixes based on confidence and entropy analysis, respectively, which are more accurate than possible fixes generated by heuristics. We experimentally verify that our techniques significantly improve the accuracy of record matching and data repairing taken as separate processes, using real-life data.

We Challenge You to Certify Your Updates

Su Chen (National University of Singapore), Xin Luna Dong (AT&T Labs-Research), Laks V.S. Lakshmanan (University of British Columbia), Divesh Srivastava (AT&T Labs-Research)

Correctness of data residing in a database is vital. While integrity constraint enforcement can often ensure data consistency, it is inadequate to protect against updates that involve careless, unintentional errors, e.g., whether a specified update to an employee's record was for the intended employee. We propose a novel approach that is complementary to existing integrity enforcement techniques, to guard against such erroneous updates. Our approach is based on (a) updaters providing an update certificate with each database update, and (b) the database system verifying the correctness of the update certificate provided before performing the update. We formalize a certificate as a (challenge, response) pair, and characterize good certificates as those that are easy for updaters to provide and, when correct, give the system enough confidence that the update was indeed intended We present algorithms that efficiently enumerate good challenges, without exhaustively exploring the search space of all challenges.We experimentally demonstrate that (i) databases have many good challenges, (ii) these challenges can be efficiently identified, (iii) certificates can be quickly verified for correctness, (iv) under natural models of an updater's knowledge of the database, update certificates catch a high percentage of the erroneous updates without imposing undue burden on the updaters performing correct updates, and (v) our techniques are robust across a wide range of challenge parameter settings.

Labeling Recursive Workflow Executions On-the-Fly

Zhuowei Bao (University of Pennsylvania), Susan B. Davidson (University of Pennsylvania), Tova Milo (Tel Aviv University)

This paper presents a compact labeling scheme for answering reachability queries over workflow executions. In contrast to previous work, our scheme allows nodes (processes and data) in the execution graph to be labeled on-the-fly, i.e., in a dynamic fashion. In this way, reachability queries can be answered as soon as the relevant data is produced. We first show that, in general, for workflows that contain recursion, dynamic labeling of executions requires long (linear-size) labels. Fortunately, most real-life scientific workflows are linear recursive, and for this natural class we show that dynamic, yet compact (logarithmic-size) labeling is possible. Moreover, our scheme labels the executions in linear time, and answers any reachability query in constant time. We also show that linear recursive workflows are, in some sense, the largest class of workflows that allow compact, dynamic labeling schemes. Interestingly, the empirical evaluation, performed over both real and synthetic workflows, shows that our proposed dynamic scheme outperforms the state-of-the-art static scheme for large executions, and creates labels that are shorter by a factor of almost 3.

Tracing Data Errors with View-Conditioned Causality

Alexandra Meliou (University of Washington), Wolfgang Gatterbauer (University of Washington), Suman Nath (Microsoft Research), Dan Suciu (University of Washington)

A surprising query result is often an indication of errors in the query or the underlying data. Recent work suggests using causal reasoning to find explanations for the surprising result. In practice, however, one often has multiple queries and/or multiple answers, some of which may be considered correct and others unexpected. In this paper, we focus on determining the causes of a set of unexpected results, possibly conditioned on some prior knowledge of the correctness of another set of results. We call this problem View-Conditioned Causality. We adapt the definitions of causality and responsibility for the case of multiple answers/views and provide a non-trivial algorithm that reduces the problem of finding causes and their responsibility to a satisfiability problem that can be solved with existing tools. We evaluate both the accuracy and

effectiveness of our approach on a real dataset of user-generated mobile device tracking data, and demonstrate that it can identify causes of error more effectively than static Boolean influence and alternative notions of causality.

SIGMOD Research 11: Information Extraction

Hybrid In-Database Inference for Declarative Information Extraction

Daisy Zhe Wang (University of California, Berkeley), Michael J. Franklin (University of California, Berkeley), Minos Garofalakis (Technical University of Crete), Joseph M. Hellerstein (University of California, Berkeley), Michael L. Wick (University of Massachusetts, Amherst)

In the database community, work on information extraction (IE) has centered on two themes: how to effectively manage IE tasks, and how to manage the uncertainties that arise in the IE process in a scalable manner. Recent work has proposed a probabilistic database (PDB) based declarative IE system that supports a leading statistical IE model, and an associated inference algorithm to answer top-k-style queries over the probabilistic IE outcome. Still, the broader problem of effectively supporting general probabilistic inference inside a PDB-based declarative IE system remains open. In this paper, we explore the in-database implementations of a wide variety of inference algorithms suited to IE, including two Markov chain Monte Carlo algorithms, the Viterbi and the sum-product algorithms. We describe the rules for choosing appropriate inference algorithms based on the model, the query and the text, considering the trade-off between accuracy and runtime. Based on these rules, we describe a hybrid approach to optimize the execution of a single probabilistic IE query to employ different inference algorithms appropriate for different records. We show that our techniques can achieve up to 10-fold speedups compared to the non-hybrid solutions proposed in the literature.

Faerie: Efficient Filtering Algorithms for Approximate Dictionary-based Entity Extraction

Guoliang Li (Tsinghua University), Dong Deng (Tsinghua University), Jianhua Feng (Tsinghua University)

Dictionary-based entity extraction identifies predefined entities (e.g., person names or locations) from a document. A recent trend for improving extraction recall is to support approximate entity extraction, which finds all substrings in the document that approximately match entities in a given dictionary. Existing methods to address this problem support either token-based similarity (e.g., Jaccard Similarity) or character-based dissimilarity (e.g., Edit Distance). It calls for a unified method to support various similarity/dissimilarity functions, since a unified method can reduce the programming efforts, hardware requirements, and the manpower. In addition, many substrings in the document have overlaps, and we have an opportunity to utilize the shared computation across the overlaps to avoid unnecessary redundant computation. In this paper, we propose a unified framework to support many similarity/dissimilarity functions, such as jaccard similarity, cosine similarity, dice similarity, edit similarity, and edit distance. We devise efficient filtering algorithms to utilize the shared computation and develop effective pruning techniques to improve the performance. The experimental results show that our method achieves high performance and outperforms state-of-the-art studies.

Joint Unsupervised Structure Discovery and Information Extraction

Eli Cortez (Universidade Federal do Amazonas), Daniel Oliveira (Universidade Federal do Amazonas), Altigran S. da Silva (Universidade Federal do Amazonas), Edleno S. de Moura (Universidade Federal do Amazonas), Alberto H. F. Laender (Universidade Federal de Minas Gerais)

In this paper we present JUDIE (Joint Unsupervised Structure Discovery and Information Extraction), a new method for automatically extracting semi-structured data records in the form of continuous text (e.g., bibliographic citations, postal addresses, classified ads, etc.) and having no explicit delimiters between them. While in state-of-the-art Information Extraction methods the structure of the data records is manually supplied the by user as a training step, JUDIE is capable of detecting the structure of each individual record being extracted without any user assistance. This is accomplished by a novel Structure Discovery algorithm that, given a sequence of labels representing attributes assigned to potential values, groups these labels into individual records by looking for frequent patterns of label repetitions among the given sequence. We also show how to integrate this algorithm in the information extraction process by means of successive refinement steps that alternate information extraction and structure discovery. Through an extensively experimental evaluation with different datasets in distinct domains, we compare JUDIE with state-of-the-art information extraction methods and conclude that, even without any user intervention, it is able to achieve high quality results on the tasks of discovering the structure of the records and extracting information from them.

Attribute Domain Discovery for Hidden Web Databases

Xin Jin (George Washington University), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

Many web databases are hidden behind restrictive form-like interfaces which may or may not provide domain information for an attribute. When attribute domains are not available, domain discovery becomes a critical challenge facing the application of a broad range of existing techniques on third-party analytical and mash-up applications over hidden databases. In this paper, we consider the problem of domain discovery over a hidden database through its web interface. We prove that for any database schema, an achievability guarantee on domain discovery can be made based solely upon the

interface design. We also develop novel techniques which provide effective guarantees on the comprehensiveness of domain discovery. We present theoretical analysis and extensive experiments to illustrate the effectiveness of our approach.

SIGMOD Research 12: Keyword Search and Ranked Queries

Keyword Search over Relational Databases: A Metadata Approach

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), Raquel Trillo Lado (University of Zaragoza,), Yannis Velegrakis (University of Trento)

Keyword queries offer a convenient alternative to traditional SQL in querying relational databases with large, often unknown, schemas and instances. The challenge in answering such queries is to discover their intended semantics, construct the SQL queries that describe them and used them to retrieve the respective tuples. Existing approaches typically rely on indices built a-priori on the database content. This seriously limits their applicability if a-priori access to the database content is not possible. Examples include the on-line databases accessed through web interface, or the sources in information integration systems that operate behind wrappers with specific query capabilities. Furthermore, existing literature has not studied to its full extend the inter-dependencies across the ways the different keywords are mapped into the database values and schema elements.In this work, we describe a novel technique for translating keyword queries into SQL based on the Munkres (a.k.a. Hungarian) algorithm. Our approach not only tackles the above two limitations, but it offers significant improvements in the identification of the semantically meaningful SQL queries that describe the intended keyword query semantics. We provide details of the technique implementation and an extensive experimental evaluation.

Sharing Work in Keyword Search over Databases

Marie Jacob (University Of Pennsylvania), Zachary Ives (University Of Pennsylvania)

An important means of allowing non-expert end-users to pose ad hoc queries whether over single databases or data integration systems is through keyword search. Given a set of keywords, the query processor finds matches across different tuples and tables. It computes and executes a set of relational sub-queries whose results are combined to produce the k highest ranking answers. Work on keyword search primarily focuses on single-database, single-query settings: each query is answered in isolation, despite possible overlap between queries posed by different users or at different times; and the number of relevant tables is assumed to be small, meaning that sub-queries can be processed without using cost-based methods to combine work. As we apply keyword search to support ad hoc data integration queries over scientific or other databases on the Web, we must reuse and combine computation. In this paper, we propose an architecture that continuously receives sets of ranked keyword queries, and seeks to reuse work across these queries. We extend multiple query optimization and continuous query techniques, and develop a new query plan scheduling module we call the ATC (based on its analogy to an air traffic controller). The ATC manages the flow of tuples among a multitude of pipelined operators, minimizing the work needed to return the top-k answers for all queries. We also develop techniques to manage the sharing and reuse of state as queries complete and input data streams are exhausted. We show the effectiveness of our techniques in handling queries over real and synthetic data sets.

Nearest Keyword Search in XML Documents

Yufei Tao (Chinese University of Hong Kong), Stavros Papadopoulos (Chinese University of Hong Kong), Cheng Sheng (Chinese University of Hong Kong), Kostas Stefanidis (Chinese University of Hong Kong)

This paper studies the {\em nearest keyword} ({\em NK}) problem on XML documents. In general, the dataset is a tree where each node is associated with one or more keywords. Given a node q and a keyword w, an NK query returns the node that is nearest to q among all the nodes associated with w. NK search is not only useful as a stand-alone operator but also as a building brick for important tasks such as XPath query evaluation and keyword search. We present an indexing scheme that answers NK queries efficiently, in terms of both practical and worst-case performance. The query cost is provably {\em logarithmic} to the number of nodes carrying the query keyword. The proposed scheme occupies space {\em linear} to the dataset size, and can be constructed by a fast algorithm. Extensive experimentation confirms our theoretical findings, and demonstrates the effectiveness of NK retrieval as a primitive operator in XML databases.

Efficient and Generic Evaluation of Ranked Queries

Wen Jin (Independent Consultant), Jignesh M Patel (University of Wisconsin--Madison)

An important feature of the existing methods for ranked top-k processing is to avoid searching all the objects in the underlying dataset, and limiting the number of random accesses to the data. However, the performance of these methods degrades rapidly as the number of random accesses increases. In this paper, we propose a novel and general sequential access scheme for top-k query evaluation, which outperforms existing methods. We extend this scheme to efficiently answer top-k queries in subspace and on dynamic data. We also study the ``dual" form of top-$k$ queries called ``ranking" queries, which returns the rank of a specified record/object, and propose an exact as well as two approximate solutions. An extensive empirical evaluation validates the robustness and efficiency of our techniques.

SIGMOD Research 13: Streams and Complex Event Processing

Changing Flights in Mid-air: A Model for Safely Modifying Continuous Queries

Kyumars Sheykh Esmaili (Systems Group), Tahmineh Sanamrad (ETH Zurich), Peter M Fischer (ETH Zurich), Nesime Tatbul (ETH Zurich)

Continuous queries can run for unpredictably long periods of time. During their lifetime, these queries may need to be adapted either due to changes in application semantics (e.g., the implementation of a new alert detection policy), or due to changes in the system's behavior (e.g., adapting performance to a changing load). While in previous works query modification has been implicitly utilized to serve specific purposes (e.g., load management), to date no research has been done that defines a general-purpose, reliable, and efficiently implementable model for modifying continuous queries at runtime. In this paper, we introduce a punctuation-based framework that can formally express arbitrary lifecycle operations on the basis of input-output mappings and basic control elements such as start or stop of queries. On top of this foundation, we derive all possible query change methods, each providing different levels of correctness guarantees and performance. We further show how these models can be efficiently realized in a state-of-the-art stream processing engine; we also provide experimental results demonstrating the key performance tradeoffs of the change methods.

How Soccer Players Would Do Stream Joins

Jens Teubner (ETH Zurich), Rene Mueller (IBM Almaden)

In spite of the omnipresence of parallel (multi-core) systems, the predominant strategy to evaluate window-based stream joins is still strictly sequential, mostly just straightforward along the definition of the operation semantics.In this work we present handshake join, a way of describing and executing window-based stream joins that is highly amenable to parallelized execution. Handshake join naturally leverages available hardware parallelism, which we demonstrate with an implementation on a modern multi-core system and on top of field-programmable gate arrays (FPGAs), an emerging technology that has shown distinctive advantages for high-throughput data processing. On the practical side, we provide a join implementation that substantially outperforms CellJoin (the fastest published result) and that will directly turn any degree of parallelism into higher throughput or larger supported window sizes. On the semantic side, our work gives a new intuition of window semantics, which we believe could inspire other stream processing algorithms or ongoing standardization efforts for stream query languages.

BE-Tree: An Index Structure to Efficiently Match Boolean Expressions over High-dimensional Discrete Space

Mohammad Sadoghi (University of Toronto), Hans-Arno Jacobsen (University of Toronto)

BE-Tree is a novel dynamic tree data structure designed to efficiently index Boolean expressions over a high-dimensional discrete space. BE-Tree copes with both high-dimensionality and expressiveness of Boolean expressions by introducing a novel two-phase space-cutting technique that specifically utilizes the discrete and finite domain properties of the space. Furthermore, BE-Tree employs self-adjustment policies to dynamically adapt the tree as the workload changes. We conduct a comprehensive evaluation to demonstrate the superiority of BE-Tree in comparison with state-of-the-art index structures designed for matching Boolean expressions.

TI: An Efficient Indexing Mechanism for Real-Time Search on Tweets

Chun Chen (Zhejiang University), Feng Li (National University of Singapore), Beng Chin Ooi (National University of Singapore), Sai Wu (National University of Singapore)

Real-time search dictates that new contents be made available for search immediately following their creation. From the database perspective, this requirement may be quite easily met by creating an up-to-date index for the contents and measuring search quality by the time gap between insertion time and availability of the index. This approach, however, poses new challenges for micro-blogging systems where thousands of concurrent users may upload their micro-blogs or tweets simultaneously. Due to the high update and query loads, conventional approaches would either fail to index the huge amount of newly created contents in real time or fall short of providing a scalable indexing service. In this paper, we propose a tweet index called the TI (Tweet Index), an adaptive indexing scheme for microblogging systems such as Twitter. The intuition of the TI is to index the tweets that may appear as a search result with high probability and delay indexing some other tweets. This strategy significantly reduces the indexing cost without compromising the quality of the search results. In the TI, we also devise a new ranking scheme by combining the relationship between the users and tweets. We group tweets into topics and update the ranking of a topic dynamically. The experiments on a real Twitter dataset confirm the efficiency of the TI.

SIGMOD Research 14: Query Processing

More Efficient Datalog Queries: Subsumptive Tabling Beats Magic Sets

K. Tuncay Tekle (LogicBlox), Yanhong A. Liu (State University of New York at Stony Brook)

Given a set of Datalog rules, facts, and a query, answers to the query can be inferred bottom-up starting with the facts or top-down starting with the query. The dominant strategies to improve the performance of answering queries are reusing answers to subqueries for top-down methods, and transforming rules based on demand from the query, such as the well-known magic sets transformation, for bottom-up methods. However, the performance of these strategies vary drastically, and the most effective method has remained unknown.This paper describes precise time and space complexity analysis for efficient implementation of Datalog queries using subsumptive tabling, a top-down evaluation method with more reuse of answers than the dominant tabling strategy, and shows that subsumptive tabling beats bottom-up evaluation of rules after magic sets transformation in both time and space complexities. It also describes subsumptive demand transformation, a novel method for transforming the rules so that bottom-up evaluation of the transformed rules mimics subsumptive tabling; we show that the time complexity of bottom-up evaluation after this transformation is equal to the the time complexity of top-down evaluation with subsumptive tabling. The paper further describes subsumption optimization, an optimization to increase the use of subsumption in subsumptive methods, and shows its application in the derivation of a well-known demand-driven pointer analysis algorithm. We support our analyses and comparisons through experiments with applications in ontology queries and program analysis.

Entangled Queries: Enabling Declarative Data-Driven Coordination

Nitin Gupta (Cornell University), Lucja Kot (Cornell University), Sudip Roy (Cornell University), Gabriel Bender (Cornell University), Johannes Gehrke (Cornell University), Christoph Koch (EPFL)

Many data-driven social and Web applications involve collaboration and coordination. The vision of declarative data-driven coordination (D3C), proposed in [9], is to support coordination in the spirit of data management: to make it data-centric and to specify it using convenient declarative languages. This paper introduces entangled queries, a language that extends SQL by constraints that allow for the coordinated choice of result tuples across queries originating from different users or applications. It is nontrivial to define a declarative coordination formalism without arriving at the general (NP-complete) Constraint Satisfaction Problem from AI. In this paper, we propose an efficiently enforcible syntactic safety condition that we argue is at the sweet spot where interesting declarative power meets applicability in large scale data management systems and applications. The key computational problem of D3C is to match entangled queries to achieve coordination. We present an efficient matching algorithm which statically analyzes query workloads and merges coordinating entangled queries into compound SQL queries. These can be sent to a standard database system and return only coordinated results. We present the overall architecture of an implemented system that contains our evaluation algorithm; we also evaluate the performance of the matching algorithm experimentally on realistic coordination workloads.

Data Generation using Declarative Constraints

Arvind Arasu (Microsoft Research), Raghav Kaushik (Microsoft Research), Jian Li (University of Maryland)

We study the problem of generating synthetic databases having declaratively specified characteristics. This problem is motivated by database system and application testing, data masking, and benchmarking. While the data generation problem has been studied before, prior approaches are either non-declarative or have fundamental limitations relating to data characteristics that they can capture and efficiently support. We argue that a natural, expressive, and declarative mechanism for specifying data characteristics is through {\em cardinality constraints}; a cardinality constraint specifies that the output of a query over the generated database have a certain cardinality. While the data generation problem is intractable in general, we present efficient algorithms that can handle a large and useful class of constraints. We include a thorough empirical evaluation illustrating that our algorithms handle complex constraints, scale well as the number of constraints increase, and outperform applicable prior techniques.

Efficient Auditing For Complex SQL queries

Raghav Kaushik (Microsoft Research), Ravi Ramamurthy (Microsoft Corporation)

We address the problem of data auditing that asks for an audit trail of all users and queries that potentially breached information about sensitive data. A lot of the previous work in data auditing has focused on providing strong privacy guarantees and studied the class of queries that can be audited efficiently while retaining the guarantees. In this paper, we approach data auditing from a different perspective. Our goal is to design an auditing system for {\em arbitrary} SQL queries containing constructs such as grouping, aggregation and correlated subqueries. Pivoted on the ability to feasibly address arbitrary queries, we study (1)~what privacy guarantees we can expect, and (2)~how we can efficiently perform auditing.

SIGMOD Research 15: Data Mining

Exact Indexing for Support Vector Machines

Hwanjo Yu (POSTECH (Pohang University of Science and Technology)), Ilhwan Ko (POSTECH (Pohang University of Science and Technology)), Youngdae Kim (POSTECH (Pohang University of Science and Technology)), Seungwon Hwang (POSTECH (Pohang University of Science and Technology)), Wook-Shin Han (Kyungpook National University)

SVM (Support Vector Machine) is a well-established machine learning methodology popularly used for classification, regression, and ranking. Recently SVM has been actively researched for rank learning and applied to various applications including search engines or relevance feedback systems. A query in such systems is the ranking function F learned by SVM. Once learning a function F or formulating the query, processing the query to find top-k results requires evaluating the entire database by F. So far, there exists no \emph{exact} indexing solution for SVM functions. Existing top-k query processing algorithms are not applicable to the machine-learned ranking functions, as they often make restrictive assumptions on the query, such as linearity or monotonicity of functions. Existing metric-based or reference-based indexing methods are also not applicable, because data points are invisible in the kernel space (SVM feature space) on which the index must be built. Existing kernel indexing methods return approximate results or fix kernel parameters. This paper proposes an exact indexing solution for SVM functions with varying kernel parameters. We first propose key geometric properties of the kernel space -- \emph{ranking instability} and \emph{ordering stability} -- which is crucial for building indices in the kernel space. Based on them, we develop an index structure \ikernel\ and processing algorithms. We then present clustering techniques in the kernel space to enhance the pruning effectiveness of the index. According to our experiments, \ikernel\ is highly effective overall producing 1$\sim$5\% of evaluation ratio on large data sets. According to our best knowledge, \ikernel\ is the first indexing solution that finds \emph{exact} top-k results of SVM functions without a full scan of data set.

Local Graph Sparsification for Scalable Clustering

Venu Satuluri (The Ohio State University), Srinivasan Parthasarathy (The Ohio State University), Yiye Ruan (The Ohio State University)

In this paper we look at how to sparsify a graph i.e. how to reduce the edgeset while keeping the nodes intact, so as to enable faster graph clustering without sacrificing quality. The main idea behind our approach is to preferentially retain the edges that are likely to be part of the same cluster. We propose to rank edges using a simple similarity-based heuristic that we efficiently compute by comparing the minhash signatures of the nodes incident to the edge. For each node, we select the top few edges to be retained in the sparsified graph. Extensive empirical results on several real networks and using four state-of-the-art graph clustering and community discovery algorithms reveal that our proposed approach realizes excellent speedups (often in the range 10-50), with little or no deterioration in the quality of the resulting clusters. In fact, for at least two of the four clustering algorithms, our sparsification consistently enables higher clustering accuracies.

Advancing Data Clustering via Projective Clustering Ensembles

Francesco Gullo (University of Calabria), Carlotta Domeniconi (George Mason University), Andrea Tagarelli (University of Calabria)

Projective Clustering Ensembles (PCE) are a very recent advance in data clustering research which combines the two powerful tools of clustering ensembles and projective clustering.Specifically, PCE enables clustering ensemble methods to handle ensembles composed by projective clustering solutions. PCE has been formalized as an optimization problem with either a two-objective or a single-objective function. Two-objective PCE has shown to generally produce more accurate clustering results than its single-objective counterpart, although it can handle the object-based and feature-based cluster representations only independently of one other. Moreover, both the early formulations of PCE do not follow any of the standard approaches of clustering ensembles, namely instance-based, cluster-based, and hybrid. In this paper, we propose an alternative formulation to the PCE problem which overcomes the above issues. We investigate the drawbacks of the early formulations of PCE and define a new single-objective formulation of the problem. This formulation is capable of treating the object- and feature-based cluster representations as a whole, essentially tying them in a distance computation between a projective clustering solution and a given ensemble. We propose two cluster-based algorithms for computing approximations to the proposed PCE formulation, which have the common merit of conforming to one of the standard approaches of clustering ensembles. Experiments on benchmark datasets have shown the significance of our PCE formulation, as both the proposed heuristics outperform existing PCE methods.

Sampling Based Algorithms for Quantile Computation in Sensor Networks

Zengfeng Huang (Hong Kong University of Science and Technology), Lu Wang (Hong Kong University of Science and Technology), Ke Yi (Hong Kong University of Science and Technology), Yunhao Liu (Hong Kong University of Science and Technology)

We study the problem of computing approximate quantiles in large-scale sensor networks communication-efficiently, a problem previously studied by Greenwald and Khana \cite{greenwald04} and Shrivastava et al.~\cite{shrivastava04:_median}. Their algorithms have a total communication cost of $O(k\log^2 n / \eps)$ and $O(k\log

u / \eps)$, respectively, where k is the number of nodes in the network, n is the total size of the data sets held by all the nodes, u is the universe size, and $\eps$ is the required approximation error. In this paper, we present a sampling based quantile computation algorithm with $O(\sqrt{kh}/\eps)$ total communication (h is the height of the routing tree), which grows sublinearly with the network size except in the pathological case $h=\Theta(k)$. In our experiments on both synthetic and real data sets, this improvement translates into a 10 to 100-fold communication reduction for achieving the same accuracy in the computed quantiles. Meanwhile, the maximum individual node communication of our algorithm is no higher than that of the previous two algorithms.

SIGMOD Research 16: Information Retrieval

Context-sensitive Ranking for Document Retrieval

Liang Jeff Chen (UC San Diego), Yannis Papakonstantinou (UC San Diego)

We study the problem of context-sensitive ranking for document retrieval, where a context is defined as a sub-collection of documents, and is specified by queries provided by domain-interested users. The motivation of context-sensitive search is that the ranking of the same keyword query generally depends on the context. The reason is that the underlying keyword statistics differ significantly from one context to another. The query evaluation challenge is the computation of keyword statistics at runtime, which involves expensive online aggregations. We appropriately leverage and extend materialized view research in order to deliver algorithms and data structures that evaluate context-sensitive queries efficiently. Specifically, a number of views are selected and materialized, each corresponding to one or more large contexts. Materialized views are used at query time to compute statistics which are used to compute ranking scores. Experimental results show that the context-sensitive ranking generally improves the ranking quality, while our materialized view-based technique improves the query efficiency.

Score-Consistent Algebraic Optimization of Full-Text Search Queries with GRAFT

Nathan Bales (UC San Diego), Alin Deutsch (UC San Diego), Vasilis Vassalos (A.U.E.B)

We address two open problems involving algebraic execution of full-text search queries. First, we show how to correctly apply traditional database rewrite optimizations to full-text algebra plans with integrated scoring, and explain why existing techniques fail. Second, we show how our techniques are applied in a generic scoring framework that supports a wide class of scoring algorithms, including algorithms seen in the literature and user-defined scoring.

Efficient Diversity-Aware Search

Albert Angel (University of Toronto), Nick Koudas (University of Toronto)

Typical approaches of ranking information in response to a user's query that return the most relevant results ignore important factors contributing to user satisfaction; for instance, the contents of a result document may be redundant given the results already examined. Motivated by emerging applications, in this work we study the problem of Diversity-Aware Search, the essence of which is ranking search results based on both their relevance, as well as their dissimilarity to other results reported. Diversity-Aware Search is generally a hard problem, and even tractable instances thereof cannot be efficiently solved by adapting existing approaches. We propose DIVGEN , an efficient algorithm for diversity-aware search, which achieves significant performance improvements via novel data access primitives. Although selecting the optimal schedule of data accesses is a hard problem, we devise the first low-overhead data access prioritization scheme with theoretical quality guarantees, and good performance in practice. A comprehensive evaluation on real and synthetic large-scale corpora demonstrates the efficiency and effectiveness of our approach.

Mining a Search Engine's Corpus: Efficient Yet Unbiased Sampling and Aggregate Estimation

Mingyang Zhang (George Washington University), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

Search engines over document corpora typically provide keyword-search interfaces. Examples include search engines over the web as well as those over enterprise and government websites. The corpus of such a search engine forms a rich source of information of analytical interest to third parties, but the only available access is by issuing search queries through its interface. To support data analytics over a search engine's corpus, one needs to address two main problems, the sampling of documents (for offline analytics) and the direct (online) estimation of aggregates, while issuing a small number of queries through the keyword-search interface. Existing work on sampling produces samples with unknown bias and may incur an extremely high query cost. Existing aggregate estimation technique suffers from a similar problem, as the estimation error and query cost can both be large for certain aggregates. We propose novel techniques which produce unbiased samples as well as unbiased aggregate estimates with small variances while incurring a query cost an order of magnitude smaller than the existing techniques. We present theoretical analysis and extensive experiments to illustrate the effectiveness of our approach.

SIGMOD Research 17: Probabilistic and Uncertain Databases

Ranking with Uncertain Scoring Functions: Semantics and Sensitivity Measures

Mohamed A Soliman (Greenplum), Ihab F Ilyas (University of Waterloo), Davide Martinenghi (Politecnico di Milano), Marco Tagliasacchi (Politecnico di Milano)

Ranking queries report the top-K results according to a user-defined scoring function. A widely used scoring function is the weighted summation of multiple scores. Often times, users cannot precisely specify the weights in such functions in order to produce the preferred order of results. Adopting uncertain/incomplete scoring functions (e.g., using weight ranges and partially-specified weight preferences) can better capture users preferences in this scenario. In this paper, we study two aspects in uncertain scoring functions. The first aspect is the semantics of ranking queries, and the second aspect is the sensitivity of computed results to refinements made by the user. We formalize and solve multiple problems under both aspects, and present novel techniques that compute query results efficiently to comply with the interactive nature of these problems.

Querying Uncertain Data with Aggregate Constraints

Mohan Yang (Shanghai Jiao Tong University), Haixun Wang (Microsoft Research Asia), Haiquan Chen (Auburn University), Wei-Shinn Ku (Auburn University)

Data uncertainty arises in many situations. A common approach to query processing uncertain data is to sample many ``possible worlds" from the uncertain data and to run queries against the possible worlds. However, sampling is not a trivial task, as a randomly sampled possible world may not satisfy known constraints imposed on the data. In this paper, we focus on an important category of constraints, the aggregate constraints. An aggregate constraint is placed on a set of records instead of on a single record, and a real-life system usually has a large number of aggregate constraints. It is a challenging task to find qualified possible worlds in this scenario, since tuple by tuple sampling is extremely inefficient because it rarely leads to a qualified possible world. In this paper, we introduce two approaches for querying uncertain data with aggregate constraints: constraint aware sampling and MCMC sampling. Our experiments show that the new approaches lead to high quality query results with reasonable cost.

Jigsaw: Efficient optimization over uncertain enterprise data

Oliver A Kennedy (EPFL), Suman Nath (Microsoft Research)

Probabilistic databases, in particular ones that allow users to externally define models or probability distributions -- so called VG-Functions -- are an ideal tool for constructing, simulating and analyzing hypothetical business scenarios. Enterprises often use such tools with parameterized models and need to explore a large parameter space in order to discover parameter values that optimize for a given goal. Parameter space is usually very large, making such exploration extremely expensive. We present Jigsaw, a probabilistic database-based simulation framework that addresses this performance problem. In Jigsaw, users define what-if style scenarios as parameterized probabilistic database queries and identify parameter values that achieve desired properties. Jigsaw uses a novel "fingerprinting" technique that efficiently identifies correlations between a query's output distribution for different parameter values. Using fingerprints, Jigsaw is able to reuse work performed for different parameter values, and obtain speedups of as much as 2 orders of magnitude for several real business scenarios.

Sensitivity Analysis and Explanations for Robust Query Evaluation in Probabilistic Databases

Bhargav Kanagal (University of Maryland), Jian Li (University of Maryland), Amol Deshpande (University of Maryland)

Probabilistic database systems have successfully established themselves as a tool for managing uncertain data. However, much of the research in this area has focused on efficient query evaluation and has largely ignored two key issues that commonly arise in uncertain data management: First, how to provide explanations for query results, e.g., Why is this tuple in my result? or Why does this output tuple have such high probability?. Second, the problem of determining the sensitive input tuples for the given query, e.g., users are interested to know the input tuples that can substantially alter the output, when their probabilities are modified (since they may be unsure about the input probability values). Existing systems provide the lineage/provenance of each of the output tuples in addition to the output probabilities, which is a boolean formula indicating the dependence of the output tuple on the input tuples. However, lineage does not immediately provide a quantitative relationship and it is not informative when we have multiple output tuples. In this paper, we propose a unified framework that can handle both the issues mentioned above to facilitate robust query processing. We formally define the notions of influence and explanations and provide algorithms to determine the top-l influential set of variables and the top-l set of explanations for a variety of queries, including conjunctive queries, probabilistic threshold queries, top-k queries and aggregation queries. Further, our framework naturally enables highly efficient incremental evaluation when input probabilities are modified (e.g., if uncertainty is resolved). Our preliminary experimental results demonstrate the benefits of our framework for performing robust query processing over probabilistic databases.

SIGMOD Research 18: Scalable Data Analytics

Fast Personalized PageRank on MapReduce

Bahman Bahmani (Stanford University), Kaushik Chakrabarti (Microsoft Research), Dong Xin (Google Inc.)

In this paper, we design a fast MapReduce algorithm for Monte Carlo approximation of personalized PageRank vectors of all the nodes in a graph. The basic idea is very efficiently doing single random walks of a given length starting at each node in the graph. More precisely, we design a MapReduce algorithm, which given a graph G and a length λ, outputs a single random walk of length λ starting at each node in G. We will show that the number of MapReduce iterations used by our algorithm is optimal among a broad family of algorithms for the problem, and its I/O efficiency is much better than the existing candidates. We will then show how we can use this algorithm to very efficiently approximate all the personalized PageRank vectors. Our empirical evaluation on real-life graph data and in production MapReduce environment shows that our algorithm is significantly more efficient than all the existing algorithms in the MapReduce setting.

Processing Theta-Joins using MapReduce

Alper Okcan (Northeastern University), Mirek Riedewald (Northeastern University)

Joins are essential for many data analysis tasks, but are not supported directly by the MapReduce paradigm. While there has been progress on equi-joins, implementation of join algorithms in MapReduce in general is not sufficiently understood. We study the problem of how to map arbitrary join conditions to Map and Reduce functions, i.e., a parallel infrastructure that controls data flow based on key-equality only. Our proposed join model simplifies creation of and reasoning about joins in MapReduce. Using this model, we derive a surprisingly simple randomized algorithm, called 1-Bucket-Theta, for implementing arbitrary joins (theta-joins) in a single MapReduce job. This algorithm only requires minimal statistics (input cardinality) and we provide evidence that for a variety of join problems, it is either close to optimal or the best possible option. For some of the problems where 1-Bucket-Theta is not the best choice, we show how to achieve better performance by exploiting additional input statistics. All algorithms can be made 'memory-aware', and they do not require any modifications to the MapReduce environment. Experiments show the effectiveness of our approach.

Llama: Leveraging Columnar Storage for Scalable Join Processing in the MapReduce Framework

Yuting Lin (National University of Singapore), Divyakant Agrawal (University of California, Santa Barbara), Chun Chen (Zhejiang University), Beng Chin Ooi (National University of Singapore), Sai Wu (National University of Singapore)

To achieve high reliability and scalability, most large-scale data warehouse systems have adopted the cluster-based architecture. In this paper, we propose the design of a new cluster-based data warehouse system, LLama, a hybrid data management system which combines the features of row-wise and column-wise database systems. In Llama, columns are formed into correlation groups to provide the basis for the vertical partitioning of tables. Llama employs a distributed file system (DFS) to disseminate data among cluster nodes. Above the DFS, a MapReduce-based query engine is supported. We design a new join algorithm to facilitate fast join processing. We present a performance study on TPC-H dataset and compare Llama with Hive, a data warehouse infrastructure built on top of Hadoop. The experiment is conducted on EC2. The results show that Llama has an excellent load performance and its query performance is significantly better than the traditional MapReduce framework based on row-wise storage.

A Platform for Scalable One-Pass Analytics using MapReduce

Boduo Li (University of Massachusetts Amherst), Edward Mazur (University of Massachusetts Amherst), Yanlei Diao (University of Massachusetts Amherst), Andrew McGregor (University of Massachusetts Amherst), Prashant Shenoy (University of Massachusetts Amherst)

Today's one-pass analytics applications tend to be data-intensive in nature and require the ability to process high volumes of data efficiently. MapReduce is a popular programming model for processing large datasets using a cluster of machines. However, the traditional MapReduce model is not well-suited for one-pass analytics, since it is geared towards batch processing and requires the data set to be fully loaded into the cluster before running analytical queries. This paper examines, from a systems standpoint, what architectural design changes are necessary to bring the benefits of the MapReduce model to incremental one-pass analytics. Our empirical and theoretical analyses of Hadoop-based MapReduce systems show that the widely-used sort-merge implementation for partitioning and parallel processing poses a fundamental barrier to incremental one-pass analytics, despite various optimizations. To address these limitations, we propose a new data analysis platform that employs hash techniques to enable fast in-memory processing, and a new frequent key based technique to extend such processing to workloads that require a large key-state space. Evaluation of our Hadoop-based prototype using real-world workloads shows that our new platform significantly improves the progress of map tasks, allows the reduce progress to keep up with the map progress, with up to 3 orders of magnitude reduction of internal data spills, and enables results to be returned continuously during the job.

SIGMOD Research 19: Graph Management

Neighborhood Based Fast Graph Search in Large Networks

Arijit Khan (University of California, Santa Barbara), Nan Li (University of California, Santa Barbara), Xifeng Yan (University of California, Santa Barbara), Ziyu Guan (University of California, Santa Barbara), Supriyo Chakraborty (University of California, Los Angeles), Shu Tao (IBM T. J. Watson Research Center)

Complex social and information network search becomes important with a variety of applications. In the core of these applications, lies a common and critical problem: Given a labeled network and a query graph, how to efficiently search the query graph in the target network. The presence of noise and the incomplete knowledge about the structure and content of the target network make it unrealistic to find an exact match. Rather, it is more appealing to find the top-k approximate matches. In this paper, we propose a neighborhood-based similarity measure that could avoid costly graph isomorphism and edit distance computation. Under this new measure, we prove that subgraph similarity search is NP hard, while graph similarity match is polynomial. By studying the principles behind this measure, we found an information propagation model that is able to convert a large network into a set of multidimensional vectors, where sophisticated indexing and similarity search algorithms are available. The proposed method, called Ness (Neighborhood Based Similarity Search), is appropriate for graphs with low automorphism and high noise, which are common in many social and information networks. Ness is not only efficient, but also robust against structural noise and information loss. Empirical results show that it can quickly and accurately find high-quality matches in large networks, with negligible cost.

A Memory Efficient Reachability Data Structure Through Bit Vector Compression

Sebastiaan J. van Schaik (University of Oxford & Utrecht University), Oege de Moor (University of Oxford)

When answering many reachability queries on a large graph, the principal challenge is to represent the transitive closure of the graph compactly, while still allowing fast membership tests on that transitive closure. Recent attempts to address this problem are complex data structures and algorithms such as Path-Tree and 3-HOP. We propose a simple alternative based on a novel form of bit-vector compression. Our starting point is the observation that when computing the transitive closure, reachable vertices tend to cluster together. We adapt the well-known scheme of word-aligned hybrid compression (WAH) to work more efficiently by introducing word partitions. We prove that the resulting scheme leads to a more compact data structure than its closest competitor, namely interval lists. In extensive and detailed experiments, this is confirmed in practice. We also demonstrate that the new technique can handle much larger graphs than alternative algorithms.

Incremental Graph Pattern Matching

Wenfei Fan (University of Edinburgh), Jianzhong Li (Harbin Institute of Technologies), Jizhou Luo (Harbin Institute of Technologies), Zijing Tan (Fudan University), Xin Wang (University of Edinburgh), Yinghui Wu (University of Edinburgh)

Graph pattern matching has become a routine process in emerging applications such as social networks. In practice a data graph is typically large, and is frequently updated with small changes. It is often prohibitively expensive to recompute matches from scratch via batch algorithms when the graph is updated. With this comes the need for incremental algorithms that compute changes to the matches in response to updates, to minimize unnecessary recomputation. This paper investigates incremental algorithms for graph pattern matching defined in terms of graph simulation, bounded simulation and subgraph isomorphism. (1) For simulation, we provide incremental algorithms for unit updates and certain graph patterns. These algorithms are optimal: in linear time in the size of the changes in the input and output, which characterizes the cost that is inherent to the problem itself. For general patterns we show that the incremental matching problem is unbounded, i.e., its cost is not determined by the size of the changes alone. (2) For bounded simulation, we show that the problem is unbounded even for unit updates and path patterns. (3) For subgraph isomorphism, we show that the problem is intractable and unbounded for unit updates and path patterns. (4) For multiple updates, we develop an incremental algorithm for each of simulation, bounded simulation and subgraph isomorphism. We experimentally verify that these incremental algorithms significantly outperform their batch counterparts in response to small changes, using real-life data and synthetic data.

Assessing and Ranking Structural Correlations in Graphs

Ziyu Guan (University of California), Jian Wu (Zhejiang University), Qing Zhang (TaoBao.com), Ambuj Singh (University of California), Xifeng Yan (University of California)

Real-life graphs not only have nodes and edges, but also have events taking place, e.g., product sales in social networks and virus infection in communication networks. Among different events, some exhibit strong correlation with the network structure, while others do not. Such structural correlation will shed light on viral influence existing in the corresponding network. Unfortunately, the traditional association mining concept is not applicable in graphs since it only works on homogeneous datasets like transactions and baskets. We propose a novel measure for assessing such structural correlations in heterogeneous graph datasets with events. The measure applies hitting time to aggregate the proximity among nodes that have the same event. In order to calculate the correlation scores for many events in a large network, we develop a scalable framework, called gScore, using sampling and approximation. By comparing to the situation where events are randomly distributed in the same network, our method is able to discover events that are highly correlated with the graph structure.

gScore is scalable and was successfully applied to the co-author DBLP network and social networks extracted from TaoBao.com, the largest online shopping network in China, with many interesting discoveries.

SIGMOD Research 20: OLAP and its Applications

Graph Cube: On Warehousing and OLAP Multidimensional Networks

Peixiang Zhao (University of Illinois at Urbana-Champaign), Xiaolei Li (Microsoft Corporation), Dong Xin (Google Inc.), Jiawei Han (University of Illinois at Urbana-Champaign)

We consider extending decision support facilities toward large sophisticated networks, upon which multidimensional attributes are associated with network entities, thereby forming the so-called multidimensional networks. Data warehouses and OLAP (Online Analytical Processing) technology have proven to be effective tools for decision support on relational data. However, they are not well-equipped to handle the new yet important multidimensional networks. In this paper, we introduce Graph Cube, a new data warehousing model that supports OLAP queries effectively on large multidimensional networks. By taking account of both attribute aggregation and structure summarization of the networks, Graph Cube goes beyond the traditional data cube model involved solely with numeric value based group-bys, thus resulting in a more insightful and structure-enriched aggregate network within every possible multidimensional space. Besides traditional cuboid queries, a new class of OLAP queries, crossboid, is introduced that is uniquely useful in multidimensional networks and has not been studied before. We implement Graph Cube by combining special characteristics of multidimensional networks with the existing well-studied data cube techniques. We perform extensive experimental studies on a series of real world data sets and Graph Cube is shown to be a powerful and efficient tool for decision support on large multidimensional networks.

MaSM: Efficient Online Updates in Data Warehouses

Manos Athanassoulis (Ecole Polytechnique Federale de Lausanne), Shimin Chen (Intel Labs), Anastasia Ailamaki (Ecole Polytechnique Federale de Lausanne), Phillip B. Gibbons (Intel Labs), Radu Stoica (Ecole Polytechnique Federale de Lausanne)

Data warehouses have been traditionally optimized for read-only query performance, allowing only offline updates at night, essentially trading off data freshness for performance. The need for 24x7 operations in global markets and the rise of online and other quickly-reacting businesses make concurrent online updates increasingly desirable. Unfortunately, state-of-the-art approaches fall short of supporting fast analysis queries over fresh data. The conventional approach of performing updates in place can dramatically slow down query performance, while prior proposals using differential updates either require large in-memory buffers or may incur significant update migration cost. This paper presents a novel approach for supporting online updates in data warehouses that overcomes the limitations of prior approaches, by making judicious use of available SSDs to cache incoming updates. We model the problem of query processing with differential updates as a type of outer join between the data residing on disks and the updates residing on SSDs. We present MaSM algorithms for performing such joins and periodic migrations, with small memory footprints, low query overhead, low SSD writes, efficient in-place migration of updates, and correct ACID support. Our experiments show that MaSM incurs only up to 7% overhead both on synthetic range scans (varying range size from 100GB to 4KB) and in a TPC-H query replay study, while also increasing the update throughput by orders of magnitude.

Latent OLAP: Data Cubes over Latent Variables

Deepak Agarwal (Yahoo! Research), Bee-Chung Chen (Yahoo! Research)

We introduce a novel class of data cube, called latent-variable cube. For many data analysis tasks, data in a database can be represented as points in a multi-dimensional space. Ordinary data cubes compute aggregate functions over these ``observed'' data points for each cell (i.e., region) in the space, where the cells have different granularities defined by hierarchies. While useful, data cubes do not provide sufficient capability for analyzing ``latent variables'' that are often of interest but not directly observed in data. For example, when analyzing users' interaction with online advertisements, observed data informs whether a user clicked an ad or not. However, the real interest is often in knowing the click probabilities of ads for different user populations. In this example, click probabilities are latent variables that are not observed but have to be estimated from data. We argue that latent variables are a useful construct for a number of OLAP application scenarios. To facilitate such analyses, we propose cubes that compute aggregate functions over latent variables. Specifically, we discuss the pitfalls of common practice in scenarios where latent variables should, but are not considered; we rigorously define latent-variable cube based on Bayesian hierarchical models and provide efficient algorithms. Through extensive experiments on both simulated and real data, we show that our method is accurate and runs orders of magnitude faster than the baseline.

E-Cube: Multi-Dimensional Event Sequence Analysis Using Hierarchical Pattern Query Sharing

Mo Liu (Worcester Polytechnic Institute), Elke Rundensteiner (Worcester Polytechnic Institute), Kara Greenfield (Worcester Polytechnic Institute), Chetan Gupta (Hewlett-Packard Labs), Song Wang (Hewlett-Packard Labs), Ismail Ari (Ozyegin University), Abhay Mehta (Hewlett-Packard Labs)

Many modern applications, including online financial feeds, tag-based mass transit systems and RFID-based supply chain management systems transmit real-time data streams. There is a need for event stream processing technology to analyze this vast amount of sequential data to enable online operational decision making. Existing techniques such as traditional online analytical processing (OLAP) systems are not designed for real-time pattern-based operations, while state-of-the-art Complex Event Processing (CEP) systems designed for sequence detection do not support OLAP operations. We propose a novel E-Cube model which combines CEP and OLAP techniques for efficient multi-dimensional event pattern analysis at different abstraction levels. Our analysis of the interrelationships in both concept abstraction and pattern refinement among queries facilitates the composition of these queries into an integrated E-Cube hierarchy. Based on this E-Cube hierarchy, strategies of drill-down (refinement from abstract to more specific patterns) and of roll-up (generalization from specific to more abstract patterns) are developed for the efficient workload evaluation. Our proposed execution strategies reuse intermediate results along both the concept and the pattern refinement relationships between queries. Based on this foundation, we design a cost-driven adaptive optimizer called Chase, that exploits the above reuse strategies for optimal E-Cube hierarchy execution. Our experimental studies comparing alternate strategies on a real world financial data stream under different workload conditions demonstrate the superiority of the Chase method. In particular, our Chase execution in many cases performs ten fold faster than the state-of-the art strategy for real stock market query workloads.

SIGMOD Research 21: Similarity Search and Queries

ATLAS: A Probabilistic Algorithm for High Dimensional Similarity Search

Jiaqi Zhai (Cornell University), Yin Lou (Cornell University), Johannes Gehrke (Cornell University)

Given a set of high dimensional binary vectors and a similarity function (such as Jaccard and Cosine), we study the problem of finding all pairs of vectors whose similarity exceeds a given threshold. The solution to this problem is a key component in many applications with feature-rich objects, such as text, images, music, videos, or social networks. In particular, there are many important emerging applications that require the use of relatively low similarity thresholds. We propose ATLAS, a probabilistic similarity search algorithm that in expectation finds a 1 - \delta fraction of all similar vector pairs. ATLAS uses truly random permutations both to filter candidate pairs of vectors and to estimate the similarity between vectors. At a 97.5% recall rate, ATLAS consistently outperforms all state-of-the-art approaches and achieves a speed-up of up to two orders of magnitude over both exact and approximate algorithms.

Flexible Aggregate Similarity Search

Yang Li (Shanghai Jiao Tong University), Feifei Li (Florida State University), Ke Yi (Hong Kong University of Science and Technology), Bin Yao (Florida State University), Min Wang (HP Labs China)

Aggregate similarity search, a.k.a. aggregate nearest neighbor (Ann) query, finds many useful applications in spatial and multimedia databases. Given a group Q of M query objects, it retrieves the most (or top-k) similar object to Q from a database P, where the similarity is an aggregation (e.g., sum, max) of the distances between the retrieved object p and all the objects in Q. In this paper, we propose an added flexibility to the query definition, where the similarity is an aggregation over the distances between p and any subset of ÆM objects in Q for some support 0 < Æ d 1. We call this new definition flexible aggregate similarity (Fann) search, which generalizes the Ann problem. Next, we present algorithms for answering Fann queries exactly and approximately. Our approximation algorithms are especially appealing, which are simple, highly efficient, and work well in both low and high dimensions. They also return near optimal answers with guaranteed constant-factor approximations in any dimensions. Extensive experiments on large real and synthetic datasets from 2 to 74 dimensions have demonstrated their superior efficiency and high quality.

Effective Data Co-Reduction for Multimedia Similarity Search

Zi Huang (The University of Queensland), Hengtao Shen (The University of Queensland), Jiajun Liu (The University of Queensland), Xiaofang Zhou (The University of Queensland)

Multimedia similarity search has been playing a critical role in many novel applications. Typically, multimedia objects are described by high-dimensional feature vectors (or points) which are organized in databases for retrieval. Although many high-dimensional indexing methods have been proposed to facilitate the search process, efficient retrieval over large, sparse and extremely high-dimensional databases remains challenging due to the continuous increases in data size and feature dimensionality. In this paper, we propose the first framework for Data Co-Reduction (DCR) on both data size and feature dimensionality. By utilizing recently developed co-clustering methods, DCR simultaneously reduces both size and dimensionality of the original data into a compact subspace, where lower bounds of the actual distances in the original space can be efficiently established to achieve fast and lossless similarity search in the lter-andrene approach. Particularly, DCR considers the duality between size and dimensionality, and achieves the optimal co-reduction, which generates the least number of candidates for actual distance computations. We conduct an extensive experimental study on large and real-life

multimedia datasets, with dimensionality ranging from 432 to 1936. Our results demonstrate that DCR outperforms existing methods significantly for lossless retrieval, especially in the presence of extremely high dimensionality.

Efficient Exact Edit Similarity Query Processing with the Asymmetric Signature Scheme

Jianbin Qin (University of New South Wales), Wei Wang (University of New South Wales), Yifei Lu (University of New South Wales), Chuan Xiao (University of New South Wales), Xuemin Lin (University of New South Wales & East Normal China University)

Given a query string Q, an edit similarity search finds all strings in a database whose edit distance with Q is no more than a given threshold t. Most existing method answering edit similarity queries rely on a signature scheme to generate candidates given the query string. We observe that the number of signatures generated by existing methods is far greater than the lower bound, and this results in high query time and index space complexities. In this paper, we show that the minimum signature size lower bound is t +1. We then propose asymmetric signature schemes that achieve this lower bound. We develop efficient query processing algorithms based on the new scheme. Several dynamic programming-based candidate pruning methods are also developed to further speed up the performance. We have conducted a comprehensive experimental study involving nine state-of-the-art algorithms. The experiment results clearly demonstrate the efficiency of our methods.

PODS PAPER ABSTRACTS

PODS Research 1: Streaming and Sampling

Get the Most Out of Your Sample: Optimal Unbiased Estimators Using Partial Information

Edith Cohen (AT&T Labs-Research), Haim Kaplan (Tel Aviv University)

Random sampling is an essential tool in the processing and transmission of data. It is used to summarize data too large to store or manipulate and meet resource constraints on bandwidth or battery power. Estimators that are applied to the sample facilitate fast approximate processing of queries posed over the original data and the value of the sample hinges on the quality of these estimators. Our work targets data sets such as request and traffic logs and sensor measurements, where data is repeatedly collected over multiple {\em instances}: time periods, locations, or snapshots. We are interested in operations, like quantiles and range, that span multiple instances. Subset-sums of these operations are used for applications ranging from planning to anomaly and change detection. Unbiased low-variance estimators are particularly effective as the relative error decreases with aggregation. The Horvitz-Thompson estimator, known to minimize variance for subset-sums over a sample of a single instance, is not optimal for multi-instance operations because it fails to exploit samples which provide partial information on the estimated quantity. We present a general principled methodology for the derivation of optimal unbiased estimators over sampled instances and aim to understand its potential. We demonstrate significant improvement in estimate accuracy of fundamental queries for common sampling schemes.

Tight Bounds for Lp Samplers, Finding Duplicates in Streams, and Related Problems

Hossein Jowhari, Mert Sağlam (Simon Fraser University), Gábor Tardos (Rényi Institute of Mathematics & Simon Fraser University)

In this paper, we present near-optimal space bounds for Lp-samplers. Given a stream of updates (additions and subtraction) to the coordinates of an underlying vector x in R^n, a perfect Lp sampler outputs the i-th coordinate with probability |x_i|^p/\|x\|_p^p. In SODA 2010, Monemizadeh and Woodruff showed polylog space upper bounds for approximate Lp-samplers and demonstrated various applications of them. Very recently, Andoni, Krauthgamer and Onak improved the upper bounds and gave a O(\epsilon^{-p}\log^3 n) space \epsilon relative error and constant failure rate Lp-sampler for p \in [1,2]. In this work, we give another such algorithm requiring only O(\epsilon^{-p}\log^2n) space for p \in (1,2). For p \in (0,1), our space bound is O(\epsilon^{-1}\log^2n), while for the p=1 case we have an O(\log(1/\epsilon)\epsilon^{-1}\log^2n) space algorithm. We also give a O(\log^2 n) bits zero relative error L_0-sampler, improving the O(\log^3 n) bits algorithm due to Frahling, Indyk and Sohler. As an application of our samplers, we give better upper bounds for the problem of finding duplicates in data streams. In case the length of the stream is longer than the alphabet size, L_1 sampling gives us an O(\log^2 n) space algorithm, thus improving the previous O(\log^3 n) bound due to Gopalan and Radhakrishnan. In the second part of our work, we prove an \Omega(\log^2 n) lower bound for sampling from 0, \pm 1 vectors (in this special case, the parameter p is not relevant for Lp sampling). This matches the space of our sampling algorithms for constant \epsilon>0. We also prove tight space lower bounds for the finding duplicates and heavy hitters problems. We obtain these lower bounds using reductions from the communication complexity problem augmented indexing.

Pan-private Algorithms Via Statistics on Sketches

Darakhshan Mir, S. Muthukrishnan, Aleksandar Nikolov, Rebecca N. Wright (Rutgers University)

Consider fully dynamic data, where we track data as it gets inserted and deleted. There are well developed notions of private data analyses with dynamic data, for example, using differential privacy. We want to go beyond privacy, and consider privacy together with security, formulated recently as pan-privacy by Dwork et al. (ICS 2010). Informally, pan-privacy preserves differential privacy while computing desired statistics on the data, even if the internal memory of the algorithm is compromised (say, by a malicious break-in or insider curiosity or by fiat by the government or law). We study pan-private algorithms for basic analyses, like estimating distinct count, moments, and heavy hitter count, with fully dynamic data. We present the first known pan-private algorithms for these problems in the fully dynamic model. Our algorithms rely on sketching techniques popular in streaming: in some cases, we add suitable noise to a previously known sketch, using a novel approach of calibrating noise to the underlying problem structure and the projection matrix of the sketch; in other cases, we maintain certain statistics on sketches; in yet others, we define novel sketches. We also present the first known lower bounds explicitly for pan privacy, showing our results to be nearly optimal for these problems. Our lower bounds are stronger than those implied by differential privacy or dynamic data streaming alone and hold even if unbounded memory and/or unbounded processing time are allowed. The lower bounds use a noisy decoding argument and exploit a connection between pan-private algorithms and data sanitization.

FIFO Indexes for Decomposable Problems

Cheng Sheng, Yufei Tao (Chinese University of Hong Kong)

This paper studies \emph{first-in-first-out} (FIFO) {\em indexes}, each of which manages a dataset where objects are deleted in the same order as their insertions. We give a technique that converts a static data structure to a FIFO index for all decomposable problems, provided that the static structure can be constructed efficiently. We present FIFO access methods to solve several problems including {\em half-plane search}, {\em nearest neighbor search}, and {\em extreme-point search}. All of our structures consume linear space, and have optimal or near-optimal query cost.

PODS Research 2: Incomplete Information and Awards

Optimal Aggregation Algorithms for Middleware.

Ronald Fagin, Amnon Lotem, and Moni Naor.

Data Exchange Beyond Complete Data

Marcelo Arenas (PUC Chile), Jorge Pérez (Universidad de Chile), Juan Reutter (University of Edinburgh)

In the traditional data exchange setting, source instances are restricted to be complete in the sense that every fact is either true or false in these instances. Although natural for a typical database translation scenario, this restriction is gradually becoming an impediment to the development of a wide range of applications that need to exchange objects that admit several interpretations. In particular, we are motivated by two specific applications that go beyond the usual data exchange scenario: exchanging incomplete information and exchanging knowledge bases.

Incomplete Information and Certain Answers in General Data Models

Leonid Libkin (University of Edinburgh)

While incomplete information is ubiquitous in all data models - especially in applications involving data translation or integration - our understanding of it is still not completely satisfactory. For example, even such a basic notion as certain answers for XML queries was only introduced recently, and in a way seemingly rather different from relational certain answers. The goal of this paper is to introduce a general approach to handling incompleteness, and to test its applicability in known data models such as relations and documents. The approach is based on representing degrees of incompleteness via semantics-based orderings on database objects. We use it to both obtain new results on incompleteness and to explain some previously observed phenomena. Specifically we show that certain answers for relational and XML queries are two instances of the same general concept; we describe structural properties behind the naive evaluation of queries; answer open questions on the existence of certain answers in the XML setting; and show that previously studied ordering-based approaches were only adequate for SQL's primitive view of nulls. We define a general setting that subsumes relations and documents to help us explain in a uniform way how to compute certain answers, and when good solutions can be found in data exchange. We also look at the complexity of common problems related to incompleteness, and generalize several results from relational and XML contexts.

Determining the Currency of Data

Wenfei Fan (University of Edinburgh & Harbin Institute of Technology), Floris Geerts (University of Edinburgh), Jef Wijsen (Université de Mons

Data in real-life databases become obsolete rapidly. One often finds that multiple values of the same entity reside in a database. While all of these values were once correct, most of them may have become stale and inaccurate. Worse still, the values often do not carry reliable timestamps. With this comes the need for studying data currency, to identify the current value of an entity in a database and to answer queries with the current values, in the absence of timestamps. This paper investigates the currency of data. (1) We propose a model that specifies partial currency orders in terms of simple constraints. The model also allows us to express what values are copied from other data sources, bearing currency orders in those sources, in terms of copy functions defined on correlated attributes. (2) We study fundamental problems for data currency, to determine whether a specification is consistent, whether a value is more current than another, and whether a query answer is certain no matter how partial currency orders are completed. (3) Moreover, we identify several problems associated with copy functions, to decide whether a copy function imports sufficient current data to answer a query, whether such a function copies redundant data, whether a copy function can be extended to import necessary current data for a query while respecting the constraints, and whether it suffices to copy data of a bounded size. (4) We establish upper and lower bounds of these problems, all matching, for combined complexity and data complexity, and for a variety of query languages. We also identify special cases that warrant lower complexity.

PODS Research 3: Index Structures and External Memory

New Results on Two-dimensional Orthogonal Range Aggregation in External Memory

Cheng Sheng, Yufei Tao (Chinese University of Hong Kong)

We consider the {\em orthogonal range aggregation} problem. The dataset S consists of N axis-parallel rectangles in $\real^2$, each of which is associated with an integer {\em weight}. Given an axis-parallel rectangle Q and an aggregate function F, a query reports the aggregated result of the weights of the rectangles in S intersecting Q. The goal is to preprocess S into a structure such that all queries can be answered efficiently.We present indexing schemes to solve the problem in external memory when $F =$ {\em max} (hence, {\em min}) and $F =$ {\em sum} (hence, {\em count} and {\em average}), respectively. Our schemes have linear or near-linear space, and answer a query in $O(\log_B N)$ or $O(\log_B^2 N)$ I/Os, where B is the disk block size.

On Finding Skylines in External Memory

Cheng Sheng , Yufei Tao (Chinese University of Hong Kong)

We consider the {\em skyline problem} (a.k.a.\ the {\em maxima problem}), which has been extensively studied in the database community. The input is a set P of d-dimensional points. A point {\em dominates} another if the former has a lower coordinate than the latter on every dimension. The goal is to find the {\em skyline}, which is the set of points $p \in P$ such that p is not dominated by any other data point. In the external-memory model, the 2-d version of the problem is known to be solvable in $O((N/B) \log_{M/B} (N/B))$ I/Os, where N is the cardinality of P, B the size of a disk block, and M the capacity of main memory. For fixed $d \ge 3$, we present an algorithm with I/O-complexity $O((N/B) \log^{d-2}_{M/B} (N/B))$. Previously, the best solution was adapted from an in-memory algorithm, and requires $O((N/B) \log^{d-2}_2 (N/M))$ I/Os.

Beyond Simple Aggregates: Indexing for Summary Queries

Zhewei Wei, Ke Yi (Hong Kong University of Science and Technology)

Database queries can be broadly classified into two categories: reporting queries and aggregation queries. The former retrieves a collection of records from the database that match the query's conditions, while the latter returns an aggregate, such as count, sum, average, or max (min), of a particular attribute of these records. Aggregation queries are especially useful in business intelligence and data analysis applications where users are interested not in the actual records, but some statistics of them. They can also be executed much more efficiently than reporting queries, by embedding properly precomputed aggregates into an index. However, reporting and aggregation queries provide only two extremes for exploring the data. Data analysts often need more insight into the data distribution than what those simple aggregates provide, and yet certainly do not want the sheer volume of data returned by reporting queries. In this paper, we design indexing techniques that allow for extracting a statistical summary of all the records in the query. The summaries we support include frequent items, quantiles, various sketches, and wavelets, all of which are of central importance in massive data analysis. Our indexes require linear space and extract a summary with the optimal or near-optimal query cost.

Space-Efficient Substring Occurrence Estimation

Alessio Orlandi (University of Pisa), Rossano Venturini (ISTI-CNR)

We study the problem of estimating the number of occurrences of substrings in textual data: A text T on some alphabet \Sigma of size \sigma is preprocessed and an index I is built. The index is used in lieu of the text to answer queries of the form Count\approx (P), returning an approximated number of the occurrences of an arbitrary pattern P as a substring of T. The problem has its main application in selectivity estimation related to the LIKE predicate in textual databases. Our focus is on obtaining an algorithmic solution with guaranteed error rates and small footprint. To achieve that, we first enrich previous work in the area of compressed text-indexing providing an optimal data structure that requires \Theta(\frac{|T| \log \sigma}{\err}) bits where \err \geq 1 is the additive error on any answer. We also approach the issue of guaranteeing exact answers for sufficiently frequent patterns, providing a data structure whose size scales with the amount of such patterns. Our theoretical findings are sustained by experiments showing the practical impact of our data structures.

PODS Research 4: Provenance

Provenance for Aggregate Queries

Yael Amsterdamer (Tel Aviv University and University of Pennsylvania), Daniel Deutch (Ben Gurion University and University of Pennsylvania), Val Tannen (University of Pennsylvania)

We study in this paper provenance information for queries with aggregation. Provenance information was studied in the context of various query languages that do not allow for aggregation, and recent work has suggested to capture provenance by annotating the different database tuples with elements of a commutative semiring and propagating the annotations through query evaluation. We show that aggregate queries pose novel challenges rendering this approach inapplicable. Consequently, we propose a new approach, where we annotate with provenance information not just tuples but also the individual values within tuples, using provenance to describe the values computation. We realize this approach in a

concrete construction, first for ``simple" queries where the aggregation operator is the last one applied, and then for arbitrary (positive) relational algebra queries with aggregation; the latter queries are shown to be more challenging in this context. Finally, we use aggregation to encode queries with difference, and study the semantics obtained for such queries on provenance annotated databases.

On Provenance Minimization

Yael Amsterdamer (Tel Aviv University and University of Pennsylvania), Daniel Deutch (Ben Gurion University and University of Pennsylvania), Tova Milo (Tel Aviv University), Val Tannen (University of Pennsylvania)

Provenance information has been proved to be very effective in capturing the computational process performed by queries, and has been used extensively as the input to many advanced data management tools (e.g. view maintenance, trust assessment, or query answering in probabilistic databases). We study here the core of provenance information, namely the part of provenance that appears in the computation of every query equivalent to the given one. This provenance core is informative as it describes the part of the computational process that is inherent to the query. It is also useful as a compact input to the above mentioned data management tools. We study algorithms that, given a query, compute an equivalent query that realizes the core provenance for all tuples in its result. We study these algorithms for queries of varying expressive power. Finally, we observe that, in general, one would not want to require database systems to evaluate a specific query that realizes the core provenance, but instead to be able to find, possibly off-line, the core provenance of a given tuple in the output (computed by an arbitrary equivalent query), without rewriting the query. We provide algorithms for such direct computation of the core provenance.

On the Complexity of Privacy-Preserving Complex Event Processing

Yeye He, Siddharth Barman (University of Wisconsin-Madison), Di Wang (Worcester Polytechnic Institute), Jeffrey F. Naughton (University of Wisconsin-Madison)

Complex Event Processing (CEP) Systems are stream processing systems that monitor incoming event streams in search of userspecified event patterns. While CEP systems have been adopted in a variety of applications, the privacy implications of event pattern reporting mechanisms have yet to be studied a stark contrast to the significant amount of attention that has been devoted to privacy for relational systems. In this paper we present a privacy problem that arises when the system must support desired patterns (those that should be reported if detected) and private patterns (those that should not be revealed). We formalize this problem, which we term privacy-preserving, utility maximizing CEP (PP-CEP), and analyze its complexity under various assumptions. Our results show that this is a rich problem to study and shed some light on the difficulty of developing algorithms that preserve utility without compromising privacy.

Provenance Views for Module Privacy

Susan B. Davidson, Sanjeev Khanna (University of Pennsylvania), Tova Milo (Tel Aviv University), Debmalya Panigrahi (Massachusetts Institute of Technology), Sudeepa Roy (University of Pennsylvania)

Scientific workflow systems increasingly store provenance information about the module executions used to produce a data item, as well as the parameter settings and intermediate data items passed between module executions. However, authors/owners of workflows may wish to keep some of this information confidential. In particular, a {\em module} may be proprietary, and users should not be able to infer its behavior by seeing mappings between all data inputs and outputs. The problem we address in this paper is the following: Given a workflow, abstractly modeled by a relation R, a privacy requirement Γ and costs associated with data. The \emph{owner} of the workflow decides which data (attributes) to hide, and provides the {\em user} with a view R' which is the projection of R over attributes which have {\em not} been hidden. {\em The goal is to minimize the cost of hidden data while guaranteeing that individual modules are Γ-private.} We call this the \secureview\ problem. We formally define the problem, study its complexity, and offer algorithmic solutions.

PODS Research 5: Queries and Views

Querying Graph Patterns

Pablo Barcelo (Universidad de Chile), Leonid Libkin, Juan L. Reutter (University of Edinburgh)

Graph data appears in a variety of application domains, and many uses of it, such as querying, matching, and transforming data, naturally result in incompletely specified graph data, i.e., graph patterns. While queries need to be posed against such data, techniques for querying patterns are generally lacking, and properties of such queries are not well understood. Our goal is to study the basics of querying graph patterns. We first identify key features of patterns, such as node and label variables and edges specified by regular expressions, and define a classification of patterns based on them. We then study standard graph queries on graph patterns, and give precise characterizations of both data and combined complexity for each class of patterns. If complexity is high, we do further analysis of features that lead to intractability, as well as lowercomplexity restrictions. We introduce a new automata model for query answering with two modes of acceptance: one captures queries returning nodes, and the other queries returning paths. We study properties of such automata, and the key

computational tasks associated with them. Finally, we provide additional restrictions for tractability, and show that some intractable cases can be naturally cast as instances of constraint satisfaction problem.

Maximizing Conjunctive Views in Deletion Propagation

Benny Kimelfeld, Jan Vondrák, Ryan Williams (IBM Research - Almaden)

In deletion propagation, tuples from the database are deleted in order to reflect the deletion of a tuple from the view. Such an operation may result in the (often necessary) deletion of additional tuples from the view, besides the intentionally deleted one. The complexity of deletion propagation is studied, where the view is defined by a conjunctive query (CQ), and the goal is to maximize the number of tuples that remain in the view. Buneman et al. showed that for some simple CQs, this problem can be solved by a trivial algorithm. This paper identifies additional cases of CQs where the trivial algorithm succeeds, and in contrast, it proves that for some other CQs the problem is NP-hard to approximate better than some constant ratio. In fact, this paper shows that among the CQs without self joins, the hard CQs are exactly the ones that the trivial algorithm fails on. In other words, for every CQ without self joins, deletion propagation is either APX-hard or solvable by the trivial algorithm. The paper then presents approximation algorithms for certain CQs where deletion propagation is APX-hard. Specifically, two constant-ratio (and polynomial-time) approximation algorithms are given for the class of star CQs without self joins. The first algorithm is a greedy algorithm, and the second is based on randomized rounding of a linear program. While the first algorithm is more efficient, the second one has a better approximation ratio. Furthermore, the second algorithm can be extended to a significant generalization of star CQs. Finally, the paper shows that self joins can have a major negative effect on the approximability of the problem.

Determining Relevance of Accesses at Runtime

Michael Benedikt, Georg Gottlob (University of Oxford), Pierre Senellart (Télécom ParisTech)

Consider the situation where a query is to be answered using Web sources that restrict the accesses that can be made on backend relational data by requiring some attributes to be given as input of the service. The accesses provide lookups on the collection of attributes values that match the binding. They can differ in whether or not they require arguments to be generated from prior accesses. Prior work has focused on the question of whether a query can be answered using a set of data sources, and in developing static access plans (e.g., Datalog programs) that implement query answering. We are interested in dynamic aspects of the query answering problem: given partial information about the data, which accesses could provide relevant data for answering a given query? We consider immediate and long-term notions of relevant accesses, and ascertain the complexity of query relevance, for both conjunctive queries and arbitrary positive queries. In the process, we relate dynamic relevance of an access to query containment under access limitations and characterize the complexity of this problem; we produce several complexity results about containment that are of interest by themselves.

Parallel Evaluation of Conjunctive Queries

Paraschos Koutris, Dan Suciu (University of Washington)

The availability of large data centers with tens of thousands of servers has led to the popular adoption of massive parallelism for data analysis on large datasets. Several query languages exist for running queries on massively parallel architectures, some based on the MapReduce infrastructure, others using proprietary implementations. Motivated by this trend, this paper analyzes the parallel complexity of conjunctive queries. We propose a very simple model of parallel computation that captures these architectures, in which the complexity parameter is the number of parallel steps requiring synchronization of all servers. We study the complexity of conjunctive queries and give a complete characterization of the queries which can be computed in one parallel step. These form a strict subset of hierarchical queries, and include flat queries like R(x,y), S(x,z), T(x,v), U(x,w), tall queries like R(x), S(x,y), T(x,y,z), U(x,y,z,w), and combinations thereof, which we call tall-flat queries. We describe an algorithm for computing in parallel any tall-flat query, and prove that any query that is not tall-flat cannot be computed in one step in this model. Finally, we present extensions of our results to queries that are not tall-flat.

PODS Research 6: Semistructured Data and XML

The Complexity of Text-Preserving XML Transformations

Timos Antonopoulos (Hasselt University and Transnational University of Limburg), Wim Martens (TU Dortmund), Frank Neven (Hasselt University and Transnational University of Limburg)

While XML is nowadays adopted as the de facto standard for data exchange, historically, its predecessor SGML was invented for describing electronic documents, i.e., marked up text. Actually, today there are still large volumes of such XML texts. We consider simple transformations which can change the internal structure of documents, that is, the mark-up, and can filter out parts of the text but do not disrupt the ordering of the words. Specifically, we focus on XML transformations where the transformed document is a subsequence of the input document when ignoring mark-up. We call the latter text-preserving XML transformations. We characterize such transformations as copy- and rearrange-free transductions. Furthermore, we study the problem of deciding whether a given XML transducer is text-preserving over a given tree language. We consider top-down transducers as well as the abstraction of XSLT called DTL. We show that deciding whether a transformation is text-preserving over an unranked regular tree language is in PTime for top-down

transducers, EXPTime-complete for DTL with XPath, and decidable for DTL with MSO patterns. Finally, we obtain that for every transducer in one of the above mentioned classes, the maximal subset of the input schema can be computed on which the transformation is text-preserving.

Efficient Evaluation for a Temporal Logic on Changing XML Documents

Mikolaj Bojańczyk (University of Warsaw), Diego Figueira (University of Warsaw & University of Edinburgh)

We consider a sequence $t_1,\ldots,t_k$ of XML documents that is produced by a sequence of local edit operations. To describe properties of such a sequence, we use a temporal logic. The logic can navigate both in time and in the document, e.g.~a formula can say that every node with label a eventually gets a descendant with label b. For every fixed formula, we provide an evaluation algorithm that works in time $O(k \cdot \log(n))$, where k is the number of edit operations and n is the maximal size of document that is produced. In the algorithm, we represent formulas of the logic by a kind of automaton, which works on sequences of documents. The algorithm works on XML documents of bounded depth.

Finding a Minimal Tree Pattern Under Neighborhood Constraints

Benny Kimelfeld (IBM Research - Almaden), Yehoshua Sagiv (The Hebrew University)

Tools that automatically generate queries are useful when schemas are hard to understand due to size or complexity. Usually, these tools find minimal tree patterns that contain a given set (or bag) of labels. The labels could be, for example, XML tags or relation names. The only restriction is that, in a tree pattern, adjacent labels must be among some specified pairs. A more expressive framework is developed here, where a schema is a mapping of each label to a collection of bags of labels. A tree pattern conforms to the schema if for all nodes v, the bag comprising the labels of the neighbors is contained in one of the bags to which the label of v is mapped. The problem at hand is to find a minimal tree pattern that conforms to the schema and contains a given bag of labels. This problem is NP-hard even when using the simplest conceivable language for describing schemas. In practice, however, the set of labels is small, so efficiency is realized by means of an algorithm that is fixed-parameter tractable (FPT). Two languages for specifying schemas are discussed. In the first, one expresses pairwise mutual exclusions between labels. Though W[1]-hardness (hence, unlikeliness of an FPT algorithm) is shown, an FPT algorithm is described for the case where the mutual exclusions form a circular-arc graph (e.g., disjoint cliques). The second language is that of regular expressions, and for that another FPT algorithm is described.

PODS Research 7: Rule-based Query Languages

A Rule-based Language for Web Data Management

Serge Abiteboul (INRIA Saclay & LSV-ENS Cachan), Meghyn Bienvenu (CNRS & Université Paris Sud), Alban Galland, Émilien Antoine (INRIA Saclay & LSV-ENS Cachan)

There is a new trend to use Datalog-style rule-based languages to specify modern distributed applications, notably on the Web. We introduce here such a language for a distributed data model where peers exchange messages (i.e. logical facts) as well as rules. The model is formally defined and its interest for distributed data management is illustrated through a variety of examples. A contribution of our work is a study of the impact on expressiveness of ``delegations" (the installation of rules by a peer in some other peer) and explicit timestamps. We also validate the semantics of our model by showing that under certain natural conditions, our semantics converges to the same semantics as the centralized system with the same rules. Indeed, we show this is even true when updates are considered.

Relational Transducers for Declarative Networking

Tom J. Ameloot, Frank Neven, Jan Van den Bussche (Hasselt University & Transnational University of Limburg)

Motivated by a recent conjecture concerning the expressiveness of declarative networking, we propose a formal computation model for eventually consistent distributed querying, based on relational transducers. A tight link has been conjectured between coordination-freeness of computations, and monotonicity of the queries expressed by such computations. Indeed, we propose a formal definition of coordination-freeness and confirm that the class of monotone queries is captured by coordination-free transducer networks. Coordination-freeness is a semantic property, but the syntactic class of obliviousstransducers we define also captures the same class of monotone queries. Transducer networks that are not coordination-free are much more powerful.

Rewrite Rules for Search Database Systems

Ronald Fagin, Benny Kimelfeld, Yunyao Li (IBM Research - Almaden), Sriram Raghavan (IBM India Research Lab), Shivakumar Vaithyanathan (IBM Research - Almaden)

The results of a search engine can be improved by consulting auxiliary data. In a search database system, the association between the user query and the auxiliary data is driven by rewrite rules that augment the user query with a set of alternative queries. This paper develops a framework that formalizes the notion of a rewrite program, which is essentially a collection of hedge-rewriting rules. When applied to a search query, the rewrite program produces a set of alternative queries that constitutes a least fixpoint (lfp). The main focus of the paper is on the lfp-convergence of a rewrite program, where a rewrite

program is lfp-convergent if the least fixpoint of every search query is finite. Determining whether a given rewrite program is lfp-convergent is undecidable; to accommodate that, the paper proposes a safety condition, and shows that safety guarantees lfp-convergence, and that safety can be decided in polynomial time. The effectiveness of the safety condition in capturing lfp-convergence is illustrated by an application to a rewrite program in an implemented system that is intended for widespread use.

SIGMOD INDUSTRIAL PAPER ABSTRACTS

SIGMOD Industrial 1: Data Management for Feeds and Streams

Bistro Data Feed Management System

Vladislav Shkapenyuk (AT&T Labs - Research), Theodore Johnson (AT&T Labs - Research), Divesh Srivastava (AT&T Labs - Research)

Data feed management is a critical component of many data intensive applications that depend on reliable data delivery to support real-time data collection, correlation and analysis. Data is typically collected from a wide variety of sources and organizations, using a range of mechanisms - some data are streamed in real time, while other data are obtained at regular intervals or collected in an ad hoc fashion. In the absence of an authoritative feed management system, individual applications are forced to make separate arrangements with feed providers, learn the structure of incoming files, monitor data quality, and trigger any processing necessary whenever new data is received. The Bistro data feed manager, designed and implemented at AT&T Labs - Research, simplifies and automates this complex task of data feed management: efficiently handling incoming raw files, identifying data feeds and distributing them to remote subscribers. A data feed management shares some similarities with publish-subscribe systems widely studied in the literature; however there are many important differences. Processing and delivering high-volume data feeds to a potentially large number of remote subscribers requires the solution of a number of technical challenges: efficient scheduling of file transmissions to maximize the cache performance, delivering data files with well-defined tardiness, feed discovery and archival, and robust handling of source and subscriber outages. In this paper we will present some of the technical solutions we developed to address these issues. Bistro supports a flexible specification language to define logical data feeds using the structure of physical data files, and to identify feed subscribers. Based on the specification, Bistro matches data files to feeds, performs file normalization and compression, efficiently delivers file, and notifies subscribers using a trigger mechanism. Data sources are often not under our management, and therefore they require careful monitoring. We describe our feed analyzer that discovers the naming structure of incoming data files to detect new feeds, dropped feeds, feed changes, or lost data in an existing feed. Bistro is currently responsible for the real-time delivery of over 100 different raw feeds within AT&T Labs, distributing data to several large-scale stream warehouses. We present performance results from these systems that demonstrate the problems that can arise, and efficacy of the Bistro implementation. We also show the kinds of feed identification problems that can arise, and how our feed analyzer allows us to rapidly identify problems or changes in data sources.

SIGMOD Industrial 2: Applying Hadoop

Apache Hadoop Goes Realtime at Facebook

Dhruba Borthakur (facebook), Jonathan Gray (facebook), Joydeep Sen Sarma (facebook), Kannan Muthukkaruppan (facebook), Nicolas Spiegelberg (facebook), Hairong Kuang (facebook), Karthik Ranganathan (facebook), Dmytro Molkov (facebook), Aravind Menon (facebook), Samuel Rash (facebook), Rodrigo Schmidt (facebook), Amitanand Aiyer (facebook)

At Facebook, we have recently deployed an application called "Facebook Messages" that is built on the Hadoop Platform. It uses Apache HBase, a database-like layer built on Hadoop. It is designed to support operating on billions of messages per day. It is the first even user-facing application that is built on the Hadoop platform. This talk describes the reasons why we chose Hadoop over other systems like Cassandra or Voldermart and discusses this applications requirement on Consistency, Availability and Partition Tolerance. We explore the enhancements we made to Hadoop to make it more real-time, the tradeoffs we made while configuring the system, and the pros-and-cons of how this is a better solution than the sharded mySQL database scheme used by mose other web-scale companies. We discuss the motivations behind our design choices, the challenges that we face in day-to-day operations and future capabilities and improvements that we are working on. We envison that this deployment is a trend-setter and new web-scale companies will prefer this Hadoop-based Hbase solution over the traditional sharded database deployments.

Nova: Continuous Pig/Hadoop Workflows

Christopher Olston (Yahoo! Research), Greg Chiou (Yahoo!), Laukik Chitnis (Yahoo!), Francis Liu (Yahoo!), Yiping Han (Yahoo!), Mattias Larsson (Yahoo!), Andreas Neumann (Yahoo!), Vellanki B. N. Rao (Yahoo!), Vijayanand Sankarasubramanian (Yahoo!), Siddharth Seth (Yahoo!), Chao Tian (Yahoo!), Topher ZiCornell (Yahoo!), Xiaodan Wang (Johns Hopkins University)

This paper describes a workflow manager developed and deployed at Yahoo called Nova, which pushes continually-arriving data through graphs of Pig programs executing on Hadoop clusters. Nova is like data stream managers in its support for stateful incremental processing, but unlike them in that it deals with data in large batches using disk-based processing. Batched incremental processing is a good fit for a large fraction of Yahoo's data processing use-cases, which deal with continually-arriving data and benefit from incremental algorithms, but do not require ultra-low-latency processing.

A Hadoop Based Distributed Loading Approach to Parallel Data Warehouses

Yu Xu (Teradata), Pekka Kostamaa (Teradata), Yan Qi (Teradata), Jian Wen (UC Riverside), Kevin Keliang Zhao (UC San Diego)

Recently the MapReduce programming paradigm, started by Google and made popular by the open source Hadoop implementation, is gaining rapid momentum in both academia and industry as another way of performing large scale data analysis. On the other hand, parallel database systems such as Teradata's parallel DBMS have been successfully deployed in large data warehouses over the last two decades for large scale business analysis in various industries over data sets ranging from a few terabytes to multiple petabytes. Large data warehouses requiring high performance continue to be built on top of powerful parallel database systems. However, due to the explosive data volume increase in recent years at some customer sites, some data such as web logs and sensor data are not managed by Teradata (EDW (Enterprise Data Warehouse), partially because it takes long time to load those extreme large volumes of data to a RDBMS. One critical part of building and running a data warehouse is the ETL (Extraction Transformation Loading) process. In fact, the growing ETL tool market is already a multi-billion-dollar market. Recent panels at various database conferences have identified that getting data into database systems has been a hindering factor to wider potential database applications such as scientific computing and social data analysis. One particular problem with the current load approaches to data warehouses is that while data are partitioned and replicated across all nodes in a parallel data warehouse, load utilities typically reside on a single node which face the issues of i) data loss/data availability if the node/hard drives crash; ii) file size limit on a single node; iii) load performance. All of these issues are mostly handled manually or only helped to some degree by some tools. We notice that one common thing between Hadoop and Teradata EDW is that data in both systems are partitioned across multiple nodes for parallel computing, which creates parallel loading opportunities not possible for DBMSs running on a single node. In this paper we describe our approach of using Hadoop as a distributed load strategy to Teradata EDW. We use Hadoop as the intermediate load server to store data to be loaded to Teradata EDW. We gain all the benefits from HDFS (Hadoop Distributed File System): i) significantly increased disk space for the file to be loaded; ii) once the data is written to HDFS, it is not necessary for the data sources to keep the data even before the file is loaded to Teradata EDW, as required in most ETL tools; iii) MapReduce programs can be used to transform and add structures to unstructured or semi-structured data; iv) more importantly since a file is distributed in HDFS, the file can be loaded more quickly in parallel to Teradata EDW, which is the main focus in our approach. We provide a fully parallel and scalable load utility called DirectLoad to efficiently load HDFS data to Teradata EDW. The key idea of the DirectLoad approach is that we first assign each data block of a HDFS file to a parallel unit in Teradata EDW, and then data blocks from Hadoop nodes are loaded directly to parallel units in Teradata EDW in parallel. Inside Teradata EDW we introduce new techniques to minimize the data movement across nodes for the DirectLoad approach. When both Hadoop and Teradata EDW coexist on the same hardware platform, as inquired by some customers because of reduced hardware and system administration costs, we have another optimization opportunity to directly load HDFS data blocks to Teradata parallel units on the same nodes. However due to the inherent non-uniform data distribution in HDFS, rarely we can avoid transferring HDFS blocks to remote Teradata nodes. We designed and implemented an optimal polynomial assignment algorithm to assign HDFS blocks to Teradata parallel units as evenly as possible and minimize remote data block tranfer to get minimal elapsed data loading time.

A Batch of PNUTS: Experiences Connecting Cloud Batch and Serving Systems

Adam E Silberstein (Yahoo! Research), Russell Sears (Yahoo! Research), Wenchao Zhou (Univeristy of Pennsylvania), Brian Frank Cooper (Yahoo! Research)

Cloud data management systems are growing in prominence, particularly at large Internet companies like Google, Yahoo!, and Amazon, which prize them for their scalability and elasticity. Each of these systems trades off between low-latency serving performance and batch processing throughput. In this paper, we discuss our experience running batch-oriented Hadoop on top of Yahoo!'s serving-oriented PNUTS system instead of the standard HDFS file system. Though PNUTS is optimized for and primarily used for serving, a number of applications at Yahoo! must run batch-oriented jobs that read or write data that is stored in PNUTS. Combining these systems reveals several key areas where the fundamental properties of each system are mismatched. We discuss the approaches we have taken to accommodate these mismatches, by either bending the batch and serving abstractions, or inventing new ones. Batch systems like Hadoop provide coarse task-level recovery, while serving systems like PNUTS provide finer record or transaction-level recovery. We combine both types to log record-level errors, while detecting and recovering from large-scale errors. Batch systems optimize for read and write throughput of large requests, while serving systems use indexing to provide low latency access to individual records. To improve latency-insensitive write throughput to PNUTS, we have created a second, batch write path. We show this write path to be competitive with batch-dedicated systems, like standard Hadoop. Batch systems provide block-level synchronous consistency among replicas, while serving systems provide record-level consistency, where asynchrony can allow records to temporarily diverge. These consistency models do not play well together, and we discuss techniques to isolate them from one another.

SIGMOD Industrial 3: Support for Business Analytics and Warehousing

Emerging Trends in the Enterprise Data Analytics

Fatma Özcan (IBM), David Hoa (IBM), Kevin S. Beyer (IBM), Andrey Balmin (IBM), Chuan Jie Liu (IBM), Yu Li (IBM)

Enterprises are dealing with ever increasing volumes of data, reaching into the petabyte scale. With many of our customer engagements, we are observing an emerging trend: They are using Hadoop-based solutions in conjunction with their data warehouses. They are using Hadoop to deal with the data volume, as well as the lack of strict structure in their data to conduct various analyses, including but not limited to Web log analysis, sophisticated data mining, machine learning and model building. This first stage of the analysis is off-line and suitable for Hadoop. But, once their data is summarized or cleansed enough, and their models are built, they are loading the results into a warehouse for interactive querying and report generation. At this later stage, they leverage the wealth of business intelligence tools, which they are accustomed to, that exist for warehouses. In this paper, we outline this use case and discuss the bidirectional connectors we developed between IBM DB2 and IBM InfoSphere BigInsights.

Efficient Processing of Data Warehousing Queries in a Split Execution Environment

Kamil Bajda-Pawlikowski (Hadapt Inc. & Yale University), Daniel J Abadi (Hadapt Inc. & Yale University), Avi Silberschatz (Yale University), Erik Paulson (University of Wisconsin-Madison)

Hadapt is a start-up company currently commercializing the Yale University research project called HadoopDB. The company focuses on building a platform for Big Data analytics in the cloud by introducing a storage layer optimized for structured data and by providing a framework for executing SQL queries efficiently. This work considers processing data warehousing queries over very large datasets. Our goal is to maximize perfor mance while, at the same time, not giving up fault tolerance and scalability. We analyze the complexity of this problem in the split execution environment of HadoopDB. Here, incoming queries are examined; parts of the query are pushed down and executed inside the higher performing database layer; and the rest of the query is processed in a more generic MapReduce framework. In this paper, we discuss in detail performance-oriented query execution strategies for data warehouse queries in split execution environments, with particular focus on join and aggregation operations. The efficiency of our techniques is demonstrated by running experiments using the TPC-H benchmark with 3TB of data. In these experiments we compare our results with a standard commercial parallel database and an open-source MapReduce implementation featuring a SQL interface (Hive). We show that HadoopDB successfully competes with other systems.

SQL Server Column Store Indexes

Per-Åke Larson (Microsoft), Cipri Clinciu (Microsoft), Eric N Hanson (Microsoft), Artem Oks (Microsoft), Susan L Price (Microsoft), Srikumar Rangarajan (Microsoft), Aleksandras Surna (Microsoft), Qingqing Zhou (Microsoft)

The SQL Server 11 release (code named Denali) introduces a new data warehouse query acceleration feature based on a new index type called a column store index. The new index type combined with new query operators processing batches of rows greatly improves data warehouse query performance: in some cases by hundreds of times and routinely a tenfold speedup for a broad range of decision support queries. Column store indexes are fully integrated with the rest of the system, including query processing and optimization. This paper gives an overview of the design and implementation of column store indexes including enhancements to query processing and query optimization to take full advantage of the new indexes. The resulting performance improvements are illustrated by a number of example queries.

An Analytic Data Engine for Visualization in Tableau

Richard Wesley (Tableau Software), Matthew Eldridge (Tableau Software), Pawel T Terlecki (Tableau Software)

Efficient data processing is critical for interactive visualization of analytic data sets. Inspired by the large amount of recent research on column-oriented stores, we have developed a new specialized analytic data engine tightly-coupled with the Tableau data visualization system. The Tableau Data Engine ships as an integral part of Tableau 6.0 and is intended for the desktop and server environments. This paper covers the main requirements of our project, system architecture and query-processing pipeline. We use real-life visualization scenarios to illustrate basic concepts and provide experimental evaluation.

SIGMOD Industrial 4: Business Analytics

LCI: A Social Channel Analysis Platform for Live Customer Intelligence

Malu Castellanos (HP), Umeshwar Dayal (HP), Meichun Hsu (HP), Riddhiman Ghosh (HP), Mohamed Dekhil (HP), Yue Lu (UIUC), Lei Zhang (University of Illinois at Chicago)

The rise of Web 2.0 with its increasingly popular social sites like Twitter, Facebook, blogs and review sites has motivated people to express their opinions publicly and more frequently than ever before. This has fueled the emerging field known as sentiment analysis whose goal is to translate the vagaries of human emotion into hard data. LCI is a social channel analysis platform that taps into what is being said to understand the sentiment with the particular ability of doing so in near real-time. LCI integrates novel algorithms for sentiment analysis and a configurable dashboard with different kinds of charts including dynamic ones that change as new data is ingested. LCI has been researched and prototyped at HP Labs in close interaction with the Business Intelligence Solutions (BIS) Division and a few customers. This paper presents an overview of the architecture and some of its key components and algorithms, focusing in particular on how LCI deals with Twitter and illustrating its capabilities with selected use cases.

SIGMOD Industrial 5: Dynamic Optimization and Unstructured Content

Turbocharging DBMS Buffer Pool Using SSDs

Jaeyoung Do (University of Wisconsin-Madison), Donghui Zhang (Microsoft Jim Gray Systems Lab), Jignesh M Patel (University of Wisconsin-Madison), David J DeWitt (Microsoft Jim Gray Systems Lab), Jeffrey F Naughton (University of Wisconsin-Madison), Alan Halverson (Microsoft Jim Gray Systems Lab)

Flash solid-state drives (SSDs) are changing the I/O landscape, which has largely been dominated by traditional hard disk drives (HDDs) for the last 50 years. In this paper we propose and systematically explore designs for using an SSD to improve the performance of a DBMS buffer manager. We propose three alternatives that differ mainly in the way that they deal with the dirty pages evicted from the buffer pool. We implemented these alternatives, as well another recently proposed algorithm for this task (TAC), in SQL Server, and ran experiments using a variety of benchmarks (TPC-C, E and H) at multiple scale factors. Our empirical evaluation shows significant performance improvements of our methods over the default HDD configuration (up to 9.4X), and up to a 6.8X speedup over TAC.

Online Reorganization in Read Optimized MMDBS

Felix Beier (Ilmenau University of Technology), Knut Stolze (IBM Research Development), Kai-Uwe Sattler (Ilmenau University of Technology)

Query performance is a critical factor in modern business intelligence and data warehouse systems. An increasing number of companies uses detailed analyses for conducting daily business and supporting management decisions. Thus, several techniques have been developed for achieving near realtime response times - techniques which try to alleviate I/O bottlenecks while increasing the throughputs of available processing units, i.e. by keeping relevant data in compressed main-memory data structures and exploiting the read-only characteristics of analytical workloads. However, update processing and skews in data distribution result in degenerations in these densely packed and highly compressed data structures affecting the memory efficiency and query performance negatively. Reorganization tasks can repair these data structures, but -- since these are usually costly operations -- require a well-considered decision which of several possible strategies should be processed and when, in order to reduce system downtimes. In this paper, we address these problems by presenting an approach for online reorganization in main-memory database systems (MMDBS). Based on a discussion of necessary reorganization strategies in IBM Smart Analytics Optimizer, a read optimized parallel MMDBS, we introduce a framework for executing arbitrary reorganization tasks online, i.e. in the background of normal user workloads without disrupting query results or performance.

Automated Partitioning Design in Parallel Database Systems

Rimma Nehme (Microsoft Jim Gray Systems Lab), Nicolas Bruno (Microsoft)

In recent years, Massively Parallel Processors (MPPs) have gained ground enabling vast amounts of data processing. In such environments, data is partitioned across multiple compute nodes, which results in dramatic performance improvements during parallel query execution. To evaluate certain relational operators in a query correctly, data sometimes needs to be re-partitioned (i.e., moved) across compute nodes. Since data movement operations are much more expensive than relational operations, it is crucial to design a suitable data partitioning strategy that minimizes the cost of such expensive data transfers. A good partitioning strategy strongly depends on how the parallel system would be used. In this paper we present a partitioning advisor that recommends the best partitioning design for an expected workload. Our tool recommends which tables should be replicated (i.e., copied into every compute node) and which ones should be distributed according to specific column(s) so that the cost of evaluating similar workloads is minimized. In contrast to previous work, our techniques are deeply integrated with the underlying parallel query optimizer, which results in more accurate recommendations in a shorter amount of time. Our experimental evaluation using a real MPP system, Microsoft SQL Server 2008 Parallel Data Warehouse, with both real and synthetic workloads shows the effectiveness of the proposed techniques and the importance of deep integration of the partitioning advisor with the underlying query optimizer.

Oracle Database Filesystem

Krishna Kunchithapadam (Oracle Corporation), Wei Zhang (Oracle Corporation), Amit Ganesh (Oracle Corporation), Niloy Mukherjee (Oracle Corporation)

Modern enterprise, web, and multimedia applications are generating unstructured content at unforeseen volumes in the form of documents, texts, and media files. Such content is generally associated with relational data such as user names, location tags, and timestamps. Storage of unstructured content in a relational database would guarantee the same robustness, transactional consistency, data integrity, data recoverability and other data management features consolidated across files and relational contents. Although database systems are preferred for relational data management, poor performance of unstructured data storage, limited data transformation functionalities, and lack of interfaces based on filesystem standards may keep more than eighty five percent of non-relational unstructured content out of databases in the coming decades. We introduce Oracle Database Filesystem (DBFS) as a consolidated solution that unifies state-of-the-art network filesystem features with relational database management ones. DBFS is a novel shared-storage network filesystem developed in the RDBMS kernel that allows content management applications to transparently store and organize files using standard filesystem interfaces, in the same database that stores associated relational content. The server component of DBFS is based on Oracle SecureFiles, a novel unstructured data storage engine within the RDBMS that provides filesystem like or better storage performance for files within the database while fully leveraging relational data management features such as transaction atomicity, isolation, read consistency, temporality, and information lifecycle management. We present a preliminary performance evaluation of DBFS that demonstrates more than 10TB/hr throughput of filesystem read and write operations consistently over a period of 12 hours on an Oracle Exadata Database cluster of four server nodes. In terms of file storage, such extreme performance is equivalent to ingestion of more than 2500 million 100KB document files a single day. The set of initial results look very promising for DBFS towards becoming the universal storage solution for both relational and unstructured content.

SIGMOD DEMONSTRATION ABSTRACTS

SIGMOD Demo A: Ranking, the Web and Social Media

SkylineSearch: Semantic Ranking and Result Visualization for PubMed

Julia Stoyanovich (University of Pennsylvania), Mayur Lodha (Columbia University), William Mee (Columbia University), Kenneth A. Ross (Columbia University)

Life sciences researchers perform scientific literature search as part of their daily activities. Many such searches are executed against PubMed, a central repository of life sciences articles, and often return hundreds, or even thousands, of results, pointing to the need for data exploration tools. In this demonstration we present SkylineSearch, a semantic ranking and result visualization system designed specifically for PubMed, and available to the scientific community at skyline.cs.columbia.edu. Our system leverages semantic annotations of articles with terms from the MeSH controlled vocabulary, and presents results as a two-dimensional skyline, plotting relevance against publication date. We demonstrate that SkylineSearch supports a richer data exploration experience than does the search functionality of PubMed, allowing users to find relevant references more easily. We also show that SkylineSearch executes queries and presents results in interactive time.

A Cross-Service Travel Engine for Trip Planning

Gang Chen (Zhejiang University), Chen Liu (National University of Singapore), Meiyu Lu (National University of Singapore), Beng Chin Ooi (National University of Singapore), Shanshan Ying (National University of Singapore), Anthony Tung (National University of Singapore), Dongxiang Zhang (National University of Singapore), Meihui Zhang (National University of Singapore)

The online travel services and resources are far from well organized and integrated. Trip planning is still a laborious job requiring interaction with a combination of services such as travel guides, personal travel blogs, map services and public transportation to piece together an itinerary. To facilitate this process, we have designed a cross-service travel engine for trip planners. Our system seamlessly and semantically integrates various types of travel services and resources based on a geographical ontology. We also built a user-friendly visualization tool for travellers to conveniently browse and design personal itineraries on Google Maps.

BRRL: A Recovery Library for Main-Memory Applications in the Cloud

Tuan Cao (Cornell University), Benjamin Sowell (Cornell University), Marcos Vaz Salles (Cornell University), Alan Demers (Cornell University), Johannes Gehrke (Cornell University)

In this demonstration we present BRRL, a library for making distributed main-memory applications fault tolerant. BRRL is optimized for cloud applications with frequent points of consistency that use data-parallelism to avoid complex concurrency control mechanisms. BRRL differs from existing recovery libraries by providing a simple table abstraction and using schema information to optimize checkpointing. We will demonstrate the utility of BRRL using a distributed transaction processing system and a platform for scientific behavioral simulations.

Tweets as Data: Demonstration of TweeQL and TwitInfo

Adam Marcus (MIT CSAIL), Michael S. Bernstein (MIT CSAIL), Osama Badar (MIT CSAIL), David R. Karger (MIT CSAIL), Samuel Madden (MIT CSAIL), Robert C. Miller (rcm@csail.mit.edu)

Microblogs such as Twitter are a tremendous repository of user-generated content. Increasingly, we see tweets used as data sources for novel applications such as disaster mapping, brand sentiment analysis, and real-time visualizations. In each scenario, the workflow for processing tweets is ad-hoc, and a lot of unnecessary work goes into repeating common data processing patterns. We introduce TweeQL, a stream query processing language that presents a SQL-like query interface for unstructured tweets to generate structured data for downstream applications. We have built several tools on top of TweeQL, most notably TwitInfo, an event timeline generation and exploration interface that summarizes events as they are discussed on Twitter. Our demonstration will allow the audience to interact with both TweeQL and TwitInfo to convey the value of data embedded in tweets.

MOBIES: Mobile-Interface Enhancement Service for Hidden Web Database

Xin Jin (George Washington University), Aditya Mone (University of Texas at Arlington), Nan Zhang (George Washington University), Gautam Das (University of Texas at Arlington)

Many web databases are hidden behind form-based interfaces which are not always easy-to-use on mobile devices because of limitations such as small screen sizes, trickier text entry, etc. In this demonstration, we have developed MOBIES, a third-party system that generates mobile-user-friendly interfaces by exploiting data analytics specific to the hidden web databases. Our user studies show the effectiveness of MOBIES on improving user experience over a hidden web database.

Search Computing: Multi-domain Search on Ranked Data

Alessandro Bozzon (Politecnico di Milano), Daniele Braga (Politecnico di Milano), Marco Brambilla (Politecnico di Milano), Stefano Ceri (Politecnico di Milano), Francesco Corcoglioniti (Politecnico di Milano), Piero Fraternali (Politecnico di Milano), Salvatore Vadacca (Politecnico di Milano)

We demonstrate the Search Computing framework for multi-domain queries upon ranked data collected from Web sources. Search Computing answers to queries like Find a good Jazz concert close to a specified location, a good restaurant and a hotel at walking distanceand fills the gap between generic and domain-specific search engines, by proposing new methods, techniques, interfaces, and tools for building search-based applications spanning multiple data services. The main enabling technology is an execution engine supporting methods for rank-join execution upon ranked data sources, abstracted and wrapped by means of a unifying service model. The demo walks through the interface for formulating multi-domain queries and follows the steps of the query engine that builds the result, with the help of run-time monitors that clearly explain the systems behavior. Once results are extracted, the demonstration shows several approaches for visualizing results and exploring the information space.

EnBlogue -- Emergent Topic Detection in Web 2.0 Streams

Foteini Alvanaki (Saarland University), Michel Sebastian (Saarland University), Krithi Ramamritham (IIT Bombay), Gerhard Weikum (Max-Planck Institute Informatics)

Emergent topics are newly arising themes in news, blogs, or tweets, often implied by interesting and unexpected correlations of tags or entities. We present the enBlogue system for emergent topic detection. The name enBlogue reflects the analogy with emerging trends in fashion often referred to as en Vogue. EnBlogue continuously monitors Web 2.0 streams and keeps track of sudden changes in tag correlations which can be adjusted using personalization to reflect particular user interests. We demonstrate enBlogue with several real-time monitoring scenarios as well as with time lapse on archived data.

NOAM: News Outlets Analysis and Monitoring System

Ilias Flaounas (University of Bristol), Omar Ali (University of Bristol), Marco Turchi (European Commission), Tristan Snowsill (University of Bristol), Florent Nicart (Université de Rouen), Tijl De Bie (University of Bristol), Nello Cristianini (University of Bristol)

We present NOAM, an integrated platform for the monitoring and analysis of news media content. NOAM is the data management system behind various applications and scientific studies aiming at modeling the media sphere. The system is also intended to address the need in the AI community for platforms where various AI technologies are integrated and deployed in the real world. It combines a relational database (DB) with state of the art AI technologies, including data mining, machine learning and natural language processing. These technologies are organized in a robust, distributed architecture of collaborating modules, that are used to populate and annotate the DB. NOAM manages tens of millions of news items in multiple languages, automatically annotating them in order to enable queries based on their semantic properties. The system also includes a unified user interface for interacting with its various modules.

SIGMOD Demo B: Systems and Performance

One-pass Data Mining Algorithms in a DBMS with UDFs

Carlos Ordonez (University of Houston), Sasi K Pitchaimalai (University of Houston)

Data mining research is extensive, but most work has proposed efficient algorithms, data structures and optimizations that work outside a DBMS, mostly on flat files. In contrast, we present a data mining system that can work on top of a relational DBMS based on a combination of SQL queries and User-Defined Functions (UDFs), debuking the common perception that SQL is inefficient or inadequate for data mining. We show our system can analyze large data sets significantly faster than external data mining tools. Moreover, our UDF-based algorithms can process a data set in one pass and have linear scalability.

Inspector Gadget: A Framework for Custom Monitoring and Debugging of Distributed Dataflows

Christopher Olston (Yahoo! Research), Benjamin Reed (Yahoo! Research)

We demonstrate a novel dataflow introspection framework called Inspector Gadget, which makes it easy to create custom monitoring and debugging add-ons to an existing dataflow engine such as Pig. The framework is motivated by a series of informal user interviews, which revealed that dataflow monitoring and debugging needs are both pressing and diverse. Of the 14 monitoring/debugging behaviors requested by users, we were able to implement 12 in Inspector Gadget, in just a few hundred lines of (Java) code each.

RAFT at Work: Speeding-Up MapReduce Applications under Task and Node Failures

Jorge-Arnulfo Quiané-Ruiz (Saarland University), Christoph Pinkel (Saarland University), Jörg Schad (Saarland University), Jens Dittrich (Saarland University)

The MapReduce framework is typically deployed on very large computing clusters where task and node failures are no longer an exception but the rule. Thus, fault-tolerance is an important aspect for the efficient operation of MapReduce jobs. However, currently MapReduce implementations fully recompute failed tasks (subparts of a job) from the beginning. This can significantly decrease the runtime performance of MapReduce applications. We present an alternative system that implements RAFT ideas. RAFT is a family of powerful and inexpensive Recovery Algorithms for Fast-Tracking MapReduce jobs under task and node failures. To recover from task failures, RAFT exploits the intermediate results persisted by MapReduce at several points in time. RAFT piggybacks checkpoints on the task progress computation. To recover from node failures, RAFT maintains a per-map task list of all input key-value pairs producing intermediate results and pushes intermediate results to reducers. In this demo, we demonstrate that RAFT recovers efficiently from both task and node failures. Further, the audience can compare RAFT with Hadoop via an easy-to-use web interface.

WattDB: An Energy-Proportional Cluster of Wimpy Nodes

Daniel Schall (TU Kaiserslautern), Volker Hudlet (TU Kaiserslautern)

The constant growth of data in all businesses leads to bigger database servers. While peak load times require fast and heavyweight hardware to guarantee performance, idle times are a waste of energy and money. Todays DBMSs have the ability to cluster several servers for performance and fault tolerance. Nevertheless, they do not support dynamic powering of the clusters nodes based on the current workload.In this demo, we propose a newly developed DBMS running on clustered commodity hardware, which is able to dynamically power nodes. The demo allows the user to interact with the DBMS and adjust workloads, while the clusters reaction is shown in real-time.

WINACS: Construction and Analysis of Web-Based Computer Science Information Networks

Tim Weninger (University of Illinois Urbana-Champaign), Marina Danilevsky (University of Illinois Urbana-Champaign), Fabio Fumarola (Universita), Joshua Hailpern (University of Illinois Urbana-Champaign), Jiawei Han (University of Illinois Urbana-Champaign), Thomas J. Johnston (University of Illinois Urbana-Champaign), Surya Kallumadi (Kansas State University), Hyungsul Kim (University of Illinois Urbana-Champaign), Zhijin Li (University of Illinois Urbana-Champaign), David McCloskey (University of Illinois Urbana-Champaign), Yizhou Sun (University of Illinois Urbana-Champaign), Nathan E. TeGrotenhuis (Whitworth University), Chi Wang (University of Illinois Urbana-Champaign), Xiao Yu (University of Illinois Urbana-Champaign)

WINACS (Web-based Information Network Analysis for Computer Science) is a project that incorporates many recent, exciting developments in data sciences to construct a Web-based computer science information network and to discover, retrieve, rank, cluster, and analyze such an information network. With the rapid development of the Web, huge amounts of information are available in the form of Web documents, structures, and links. It has been a dream of the database and Web communities to harvest such information and reconcile the unstructured nature of the Web with the neat, semi-structured schemas of the database paradigm. Taking computer science as a dedicated domain, WINACS first discovers related Web entity structures, and then constructs a heterogeneous computer science information network in order to rank, cluster and analyze this network and support intelligent and analytical queries.

A Data-oriented Transaction Execution Engine and Supporting Tools

Ippokratis Pandis (Carnegie Mellon University), Pinar Tozun (Ecole Polytechnique Federale de Lausanne), Miguel Branco (Ecole Polytechnique Federale de Lausanne), Dimitris Karampinas (University of Patras), Danica Porobic (Ecole Polytechnique Federale de Lausanne), Ryan Johnson (University of Toronto), Anastasia Ailamaki (Ecole Polytechnique Federale de Lausanne)

Conventional OLTP systems assign each transaction to a worker thread and that thread accesses data, depending on what the transaction dictates. This thread-to-transaction work assignment policy leads to unpredictable accesses. The unpredictability forces each thread to enter a large number of critical sections for the completion of even the simplest of the transactions; leading to poor performance and scalability on modern manycore hardware. This demonstration highlights the chaotic access patterns of conventional OLTP designs which are the source of scalability problems. Then, it presents a working prototype of a transaction processing engine that follows a non-conventional architecture, called data-oriented or DORA. DORA is designed around the thread-to-data work assignment policy. It distributes the transaction execution to multiple threads and offers predictable accesses. By design, DORA can decentralize the lock management service, and thereby eliminate the critical sections executed inside the lock manager. We explain the design of the system and show that it more efficiently utilizes the abundant processing power of modern hardware, always contrasting it against the conventional execution. In addition, we present different components of the system, such as a dynamic load balancer. Finally, we present a set of tools that enable the development of applications that use DORA.

iGraph in Action: Performance Analysis of Disk-Based Graph Indexing Techniques

Wook-Shin Han (Kyungpook National University), Minh-Duc Pham (Kyungpook National University), Jinsoo Lee (Kyungpook National University), Romans Kasperovics (Kyungpook National University), Jeffrey Xu Yu (Chinese University of Hong Kong)

Graphs provide a powerful way to model complex structures such as chemical compounds, proteins, images, and program dependence. The previous 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 observed that this practice may result in several problems [6]. In order to address these problems, we have implemented all representative graph indexing techniques on a common framework called iGraph [6]. In this demonstration we showcase iGraph and its visual tools using several real datasets and their workloads. For selected queries of the workloads, we show several unique features including visual performance analysis.

StreamRec: A Real-Time Recommender System

Badrish Chandramouli (Microsoft Research), Justin J Levandoski (University of Minnesota), Ahmed Eldawy (University of Minnesota), Mohamed F Mokbel (University of Minnesota)

This demonstration proposes StreamRec, a novel approach to building recommender systems that leverages a stream processing system to implement the end-to-end recommendation process. We demonstrate how the popular collaborative filtering recommendation method can be implemented as a stream-based continuous query plan. We also provide an application scenario using StreamRec as the underlying engine behind two applications: (1) a social news and (2) a movie recommendation application. StreamRec is implemented using the Microsoft StreamInsight stream processing system.

SIGMOD Demo C: Data Integration and Probabilistic Databases

Pay-As-You-Go Mapping Selection in Dataspaces

Cornelia Hedeler (The University of Manchester), Khalid Belhajjame (The University of Manchester), Norman W Paton (The University of Manchester), Alvaro A A Fernandes (The University of Manchester), Suzanne M Embury (The University of Manchester), Lu Mao (The University of Manchester), Chenjuan Guo (The University of Manchester)

The vision of dataspaces proposes an alternative to classical data integration approaches with reduced up-front costs followed by incremental improvement on a pay-as-you-go basis. In this paper, we demonstrate DSToolkit, a system that allows users to provide feedback on results of queries posed over an integration schema. Such feedback is then used to annotate the mappings with their respective precision and recall. The system then allows a user to state the expected levels of precision (or recall) that the query results should exhibit and, in order to produce those results, the system selects those mappings that are predicted to meet the stated constraints.

Exelixis: Evolving Ontology-Based Data Integration System

Haridimos Kondylakis (FORTH-ICS), Dimitris Plexousakis (FORTH-ICS)

The evolution of ontologies is an undisputed necessity in ontology-based data integration. Yet, few research efforts have focused on addressing the need to reflect ontology evolution onto the underlying data integration systems. We present Exelixis, a web platform that enables query answering over evolving ontologies without mapping redefinition. This is achieved by rewriting queries among ontology versions. First, changes between ontologies are automatically detected and described using a high level language of changes. Those changes are interpreted as sound global-as-view (GAV) mappings. Then query expansion is applied in order to consider constraints from the ontology and unfolding to apply the GAV mappings. Whenever equivalent rewritings cannot be produced we a) guide query redefinition and/or b) provide the best over-approximations, i.e. the minimally-containing and minimally-generalized rewritings. For the demonstration we will use four versions of the CIDOC-CRM ontology and real user queries to show the functionality of the system. Then we will allow conference participants to directly interact with the system to test its capabilities.

U-MAP: A System for Usage-Based Schema Matching and Mapping

Hazem Elmeleegy (AT&T Labs - Research), Jaewoo Lee (Purdue University), El Kindi Rezig (Purdue University), Mourad Ouzzani (Purdue University), Ahmed Elmagarmid (Qatar Computing Research Institute, Qatar Foundation)

This demo shows how usage information buried in query logs can play a central role in data integration and data exchange. More specifically, our system U-Map uses query logs to generate correspondences between the attributes of two different schemas and the complex mapping rules to transform and restructure data records from one of these schemas to another. We introduce several novel features showing the benefit of incorporating query log analysis into these key components of data integration and data exchange systems.

The SystemT IDE: An Integrated Development Environment for Information Extraction Rules

Laura Chiticariu (IBM Research - Almaden), Vivian Chu (IBM research - Almaden), Sajib Dasgupta (IBM Research - Almaden), Thilo W Goetz (IBM Software - Germany), Howard Ho (IBM Research - Almaden), Rajasekar Krishnamurthy (IBM Research - Almaden), Alexander Lang (IBM Software - Germany), Yunyao Li (IBM Research - Almaden), Bin Liu (University of Michigan), Sriram Raghavan (IBM Research - India), Frederick R Reiss (IBM Research - Almaden), Shivakumar Vaithyanathan (IBM Research - Almaden), Huaiyu Zhu (IBM Research - Almaden)

Information Extraction (IE)the problem of extracting structured information from unstructured text has become the key enabler for many enterprise applications such as semantic search, business analytics and regulatory compliance. While rule-based IE systems are widely used in practice due to their well-known explainability, developing high-quality information extraction rules is known to be a labor-intensive and time-consuming iterative process.Our demonstration showcases SystemT IDE, the integrated development environment for SystemT, a state-of-the-art rule based IE system from IBMResearch that has been successfully embedded in multiple IBM enterprise products. SystemT IDE facilitates the development, test and analysis of high-quality IE rules by means of sophisticated techniques, ranging from data management to machine learning. We show how to build high-quality IE annotators using a suite of tools provided by SystemT IDE, including computing data provenance, learning basic features such as regular expressions and dictionaries, and automatically refining rules based on labeled examples.

ProApproX: A Lightweight Approximation Query Processor over Probabilistic Trees

Pierre Senellart (Institut Télécom; Télécom ParisTech), Asma Souihli (Institut Télécom; Télécom ParisTech)

We demonstrate a system for querying probabilistic XML documents with simple XPath queries. A user chooses between a variety of query answering techniques, both exact and approximate, and observes the running behavior, pros, and cons, of each method, in terms of efficiency, precision of the result, and data model and query language supported.

SPROUT²: A Squared Query Engine for Uncertain Web Data

Robert Fink (University of Oxford), Andrew Hogue (Google Inc.), Dan Olteanu (University of Oxford), Swaroop Rath (University of Oxford)

SPROUT² is a query answering system that allows users to ask structured queries over tables embedded in Web pages, over Google Fusion tables, and over uncertain tables that can be extracted from answers to Google Squared. At the core of this service lies SPROUT, a query engine for probabilistic databases. This demonstration allows users to compose and ask ad-hoc queries of their choice and also to take a tour through the system's capabilities along pre-arranged scenarios on, e.g., movie actors and directors, biomass facilities, or leveraging corporate databases.

Fuzzy Prophet: Parameter exploration in uncertain enterprise scenarios

Oliver A Kennedy (EPFL), Steve Lee (Microsoft Corporation), Charles Loboz (Microsoft Corporation), Slawek Smyl (Microsoft Corporation), Suman Nath (Microsoft Research)

We present Fuzzy Prophet, a probabilistic database tool for constructing, simulating and analyzing business scenarios with uncertain data. Fuzzy Prophet takes externally defined probability distribution (so called VG-Functions) and a declarative description of a target scenario, and performs Monte Carlo simulation to compute probability distribution of the scenario's outcomes. In addition, Fuzzy Prophet supports parameter optimization,where probabilistic models are parameterized and a large parameter space must be explored to find parameters that optimize or achieve a desired goal. Fuzzy Prophet's key innovation is to use 'fingerprints' that can identify parameter values producing correlated outputs of a user-provided stochastic function and to reuse computations across such values. Fingerprints significantly expedite the process of parameter exploration in offline optimization and interactive what-if exploration tasks.

LinkDB: A Probabilistic Linkage Database System

Ekaterini Ioannou (Technical University of Crete), Wolfgang Nejdl (L3S Research Center), Claudia Niederée (L3S Research Center), Yannis Velegrakis (University of Trento)

Entity linkage deals with the problem of identifying whether two pieces of information represent the same real world object. The traditional methodology computes the similarity among the entities, and then merges those with similarity above some specific threshold. We demonstrate LinkDB, an original entity storage and querying system that deals with the entity linkage problem in a novel way. LinkDB is a probabilistic linkage database that uses existing linkage techniques to generate linkages among entities, but instead of performing the merges based on these linkages, it stores them alongside the data and performs only the required merges at run-time, by effectively taking into consideration the query specifications. We explain the technical challenges behind this kind of query answering, and we show how this new mechanism is able to provide answers that traditional entity linkage mechanisms cannot.

SIGMOD Demo D: User Support and Development Environments

CONFLuEnCE: CONtinuous workFLow ExeCution Engine

Panayiotis Neophytou (University of Pittsburgh), Panos K. Chrysanthis (University of Pittsburgh), Alexandros Labrinidis (University of Pittsburgh)

Traditional workflow enactment systems view a workflow as a one-time interaction with various data sources, executing a series of steps once, whenever the workflow results are requested. The fundamental underlying assumption has been that data sources are passive and all interactions are structured along the request/reply (query) model. Hence, traditional Workflow Management Systems cannot effectively support business or scientific reactive applications that require the processing of continuous data streams. In this demo, we will present our prototype which transforms workflow execution from the traditional step-wise workflow execution model to a continuous execution model, in order to handle data streams published and delivered asynchronously from multiple sources. We will demonstrate a supply chain management scenario which takes advantage of our continuous execution model to enable on-line interaction between different user roles as well as streaming data coming from various sources.

Demonstration of Qurk: A Query Processor for HumanOperators

Adam Marcus (MIT CSAIL), Eugene Wu (MIT CSAIL), David R. Karger (MIT CSAIL), Samuel Madden (MIT CSAIL), Robert C. Miller (MIT CSAIL)

Crowdsourcing technologies such as Amazon's Mechanical Turk (``MTurk'') service have exploded in popularity in recent years. These services are increasingly used for complex human-reliant data processing tasks, such as labelling a collection of images, combining two sets of images to identify people that appear in both, or extracting sentiment from a corpus of text snippets. There are several challenges in designing a workflow that filters, aggregates, sorts and joins human-generated data sources. Currently, crowdsourcing-based workflows are hand-built, resulting in increasingly complex programs. Additionally, developers must hand-optimize tradeoffs among monetary cost, accuracy, and time to completion of results. These challenges are well-suited to a declarative query interface that allows developers to describe their workflow at a high level and automatically optimizes workflow and tuning parameters. In this demonstration, we will present Qurk, a novel query system that allows human-based processing for relational databases. The audience will interact with the system to build queries and monitor their progress. The audience will also see Qurk from an MTurk user's perspective, and complete several tasks to better understand how a query is processed.

Automatic Example Queries for Ad Hoc Databases

Bill Howe (University of Washington), Garret Cole (University of Washington), Nodira Khoussainova (University of Washington), Leilani Battle (University of Washington)

Motivated by eScience applications, we explore automatic generation of example ""starter"" queries over unstructured collections of tables without relying on a schema, a query log, or prior input from users. Such example queries are demonstrably sufficient to have non-experts self-train and become productive using SQL, helping to increase the uptake of database technology among scientists. Our method is to learn a model for each relational operator based on example queries from public databases, then assemble queries syntactically operator-by-operator. For example, the likelihood that a pair of attributes will be used as a join condition in an example query depends on the cardinality of their intersection, among other features. Our demonstration illustrates that datasets with different statistical properties lead to different sets of example queries with different properties.

NetTrails: A Declarative Platform for Maintaining and Querying Provenance in Distributed Systems

Wenchao Zhou (University of Pennsylvania), Qiong Fei (University of Pennsylvania), Shengzhi Sun (University of Pennsylvania), Tao Tao (University of Pennsylvania), Andreas Haeberlen (University of Pennsylvania), Zachary Ives (University of Pennsylvania), Boon Thau Loo (University of Pennsylvania), Micah Sherr (Georgetown University)

We demonstrate NetTrails, a declarative platform for maintaining and interactively querying network provenance in a distributed system. Network provenance describes the history and derivations of network state that result from the execution of a distributed protocol. It has broad applicability in the management, diagnosis, and security analysis of networks. Our demonstration shows the use of NetTrails for maintaining and querying network provenance in a variety of distributed settings, ranging from declarative networks to unmodified legacy distributed systems. We conclude our demonstration with a discussion of our ongoing research on enhancing the query language and security guarantees.

GBLENDER: Visual Subgraph Query Formulation Meets Query Processing

Changjiu Jin (Nanyang Technological University), Sourav S Bhowmick (Nanyang Technological University), Xiaokui Xiao (Nanyang Technological University), Byron Choi (Hong Kong Baptist University), Shuigeng Zhou (Fudan University)

Due to the complexity of graph query languages, the need for visual query interfaces that can reduce the burden of query formulation is fundamental to the spreading of graph data management tools to wider community. We present a novel HCI (human-computer interaction)-aware graph query processing paradigm, where instead of processing a query graph after its

construction, it interleaves visual query construction and processing to improve system response time. We demonstrate a system called GBLENDER that exploits GUI latency to prune false results and prefetch candidate data graphs by employing a novel action-aware indexing scheme and a data structure called spindle-shaped graphs (SPIG). We demonstrate various innovative features of GBLENDER and its promising performance in evaluating subgraph containment and similarity queries.

Coordination through Querying in the Youtopia System

Nitin Gupta (Cornell University), Lucja Kot (Cornell University), Gabriel Bender (Cornell University), Sudip Roy (Cornell University), Johannes Gehrke (Cornell University), Christoph Koch (EPFL)

In a previous paper, we laid out the vision of declarative data-driven coordination (D3C) where users are provided with novel abstractions that enable them to communicate and coordinate through declarative specifications. In this demo, we will show Youtopia, a novel database system which is our first attempt at implementing this vision. Youtopia provides coordination abstractions within the DBMS. Users submit queries that come with explicit coordination constraints to be met by other queries in the system. Such queries are evaluated together; the system ensures that their joint execution results in the satisfaction of all coordination constraints. That is, the queries coordinate their answers in the manner specified by the users.We show how Youtopia and its abstractions simplify the implementation of a three-tier flight reservation application that allows users to coordinate travel arrangements with their friends.

DBWiki: A Structured Wiki for Curated Data and Collaborative Data Management

Peter Buneman (University of Edinburgh), James Cheney (University of Edinburgh), Sam Lindley (University of Edinburgh), Heiko Müller (CSIRO)

Wikis have proved enormously successful as a means to collaborate in the creation and publication of textual information. At the same time, a large number of curated databases have been developed through collaboration for the dissemination of structured data in specific domains, particularly bioinformatics. We demonstrate a general-purpose platform for collaborative data management, DBWiki, designed to achieve the best of both worlds. Our system not only facilitates the collaborative creation of a database; it also provides features not usually provided by database technology such as versioning, provenance tracking, citability, and annotation. In our demonstration we will show how DBWiki makes it easy to create, correct, discuss and query structured data, placing more power in the hands of users while managing tedious details of data curation automatically.

Rapid Development of Web-Based Query Interfacesfor XML Datasets with QURSED

Abhijith Kashyap (SUNY at Buffalo), Michalis Petropoulos (SUNY at Buffalo)

We present QURSED, a system that automates the development of web-based query forms and reports (QFRs) for semi-structured XML data. Whereas many tools for automated development of QFRs have been proposed for relational datasets, QURSED- to the best of our knowledge, is the first tool that facilitates development of web-based QFRs over XML data. The QURSED system is available online at http://db.cse.buffalo.edu/qursed.

UNDERGRADUATE POSTER COMPETITION

Study of techniques for the indirect enhancement of training data in the training of neural networks used in re-ranking of web search results
Panagiotis Parchas, School of Electrical and Computer Engineering, Department of Computer Science, National Technical University of Athens (NTUA), Greece

ATLAS: A Probabilistic Algorithm for High Dimensional Similarity Search
Jiaqi Zhai, Cornell University, USA

Outsourced Computation Verification
Roy Luo, University of California, Berkeley, USA

DYNAMO - Dynamic Web-Service Mashups
Apostolos Nydriotis, ECE Department, Technical University of Crete, Greece

DEMA: Dynamic Clustering of Spatio-Temporal Dataset to Improve Indexing Performance
Wook-Hee Kim, Electrical and Computer Engineering, Ulsan National Institute of Science and Technology, Korea

Link Prediction in Annotation Graphs using Graph Summarization
Philip Anderson, Computer Science Department, University of Maryland, USA

GRADUATE POSTER ABSTRACTS

Research on Moving Objects with Multimodal Transportation Modes

Jianqiu Xu (FernUniversität Hagen)

In spite of massive research having been conducted on moving objects databases, existing methods can not manage moving objects with transportation modes and covering multiple environments, e.g., Walk → Car → Indoor. The current techniques target a single environment, for example road network. But humans' movement can include several environments like roads, pavement areas, buildings, which imposes new challenges regarding representation and management of these data and efficient query processing in a database system. In this paper, we first present a data model for representing generic moving objects in multiple environments. The method is to let the space where moving objects are located consist of a set of called infrastructure objects like streets, pavements, moving buses, and then the location of a moving object is represented by referencing to these infrastructure objects. In the meantime, transportation modes are seamlessly integrated. Second, to evaluate the performance of a database system, a benchmark comprising of a scalable dataset and a set of queries is required. We propose a method to create all infrastructure objects as well as generic moving objects due to the major challenge of getting the real data. A list of queries is given which will be used for the evaluation.

Database Forensics in the Service of Information Accountability

Kyriacos E. Pavlou (The University of Arizona)

Regulations and societal expectations have recently expressed the need to mediate access to valuable databases, even by insiders. At one end of the spectrum is the approach of re- stricting access to information and on the other that of information accountability. The focus of the proposed work is effecting information accountability of data stored in databases. One way to ensure appropriate use and thus end-to-end accountability of such information is tamper detection in databases via a continuous assurance technology based on cryptographic hashing. In our current research we are working to show how to develop the necessary approaches and ideas to support accountability in high-performance databases. This will include the design of a reference architecture for information accountability and several of its variants, the development of forensic analysis algorithms and their cost model, and a systematic formulation of forensic analysis for determining when the tampering occurred and what data were tampered with. Finally, for privacy, we would like to create mechanisms for allowing as well as (temporarily) preventing the physical deletion of records in a monitored database. In order to evaluate our ideas we will design and implement an integrated tamper detection and forensic analysis system. This work will show that information accountability is a viable alternative to information restriction for ensuring the correct storage, use, and maintenance of databases.

PROGRAMMING CONTEST FINALISTS

Team 1

Jung-Sang Ahn - Korea Advanced Institute of Science and Technology

Team 2

Leif Walsh (Stony Brook University)

Team 3

Ehab Abdelhamid (King Abdullah University of Science and Technology), Majed Sahli (King Abdullah University of Science and Technology), Razen Alharbi (King Abdullah University of Science and Technology)

Team 4

Thomas Kissinger (Technische Universität Dresden), Benjamin Schlegel (Technische Universität Dresden)

Team 5

Ahmed Eldawy (University of Minnesota), Emery Mizero (University of Minnesota), Mohamed E. Khalefa (University of Minnesota)

CO-LOCATED WORKSHOPS

10th ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE '11)

http://www.itee.uq.edu.au/~mobide11

This is the tenth of a successful series of workshops that aims to actas a bridge between the data management, wireless networking, and mobile computing communities. The workshop will serve as a forum for researchers and technologists to discuss the state-of-the-art, present their contributions, and set future directions in data management for mobile and wireless access.

14th International Workshop on the Web and Databases (WebDB 2011)

http://webdb2011.rutgers.edu/

The WebDB workshop has been held thirteen times so far: the first WebDB workshop was collocated with EDBT 1998, whereas the other twelve WebDB workshop editions were collocated with the annual SIGMOD/PODS conference. The WebDB workshop provides a forum where researchers, theoreticians, and practitioners can share their insights and their knowledge on problems and solutions at the intersection of data management and the Web. WebDB has high impact and has been a forum in which a number of seminal papers have been presented. Papers on all aspects of the Web and Databases are solicited. To provide focus, since 2002 each WebDB workshop has a theme. The theme for 2011 is User Generated Content on the Web: theory, techniques, applications, systems, and people.

6th International Workshop on Networking Meets Databases (NetDB 2011)

http://research.microsoft.com/en-us/UM/people/srikanth/netdb11/

The International Workshop on Networking Meets Databases (NetDB 2011) will bring together researchers from the systems and networking and databases communities. Many research areas, such as cloud computing, data center networking, privacy-aware systems, sensor networks, network management, P2P systems, rule mining, inference over system logs and network traffic data, and declarative system-building, are blurring the boundaries between these two communities. The goal of the workshop is to foster an environment in which researchers from both communities can discuss ideas that will shape and influence these emerging research areas. We encourage submissions of early work, with novel and interesting ideas. We expect that work introduced at NetDB 2011, once fully thought through, completed, and described in a finished form, may be relevant to conferences such as SOSP, OSDI, SIGCOMM, SIGMOD, VLDB, NSDI, or ICDE.

1st ACM Workshop on Databases and Social Networks (DBSocial)

https://sites.google.com/site/dbsocial11/call-for-papers

The First ACM Workshop on Databases and Social Networks (DBSocial), to be held in conjunction with SIGMOD 2011, in Athens, Greece, is a venue for database research applied to the problems of extraction, querying, and analysis of social networks, and for database research and practice that advances the state-of-the-art in the observation, management, and analysis of inherently networked data originating primarily from social phenomena. DBSocial welcomes both theoretical and practical papers whose approaches are within the scope of databases and very closely related areas (e.g., data mining and information retrieval), and whose validation is on par with the high standards in the database community.

3th International Workshop on Semantic Web Information Management - SWIM 2011

http://mais.dia.uniroma3.it/SWIM2011/Home.html

As the World Wide Web grows, it is becoming more and more complex for humans to efficiently find and exploit the information we need. The underlying idea of having a description of data available on the Web, organized in such a way as to be used by machines for automation, integration and reuse across various applications, has been promoted by a number of research fields. The Third International Workshop on "*Semantic Web Information Management*" (SWIM) aims to review the most recent data-centered solutions for the Semantic Web. In particular, the workshop's ambition is to present and analyze techniques for semantic information management, taking advantage of the synergies between the logical basis of the Semantic Web and the logical foundations of conceptual modeling. Indeed, a *leitmotif* of these research areas is the proposal of models and methods to represent and manage appropriately structured data, permitting it to be easily machine-processable on the Web. The long-standing experience of the information modeling community can provide a significant contribution to the substantial problems arising in semantic data management using technologies such as RDF, RDFS and OWL.

7th International Workshop on Data Management on New Hardware (DaMoN 2011)
http://www.cse.ust.hk/damon2011/

The aim of this one-day workshop is to bring together researchers who are interested in optimizing database performance on modern computing infrastructure by designing new data management techniques and tools. The continued evolution of computing hardware and infrastructure imposes new challenges and bottlenecks to program performance. As a result, traditional database architectures that focus solely on I/O optimization increasingly fail to utilize hardware resources efficiently. CPUs with superscalar out-of-order execution, simultaneous multi-threading, multi-level memory hierarchies, and future storage hardware (such as MEMS) impose a great challenge to optimizing database performance. Consequently, exploiting the characteristics of modern hardware has become an important topic of database systems research. The goal is to make database systems adapt automatically to the sophisticated hardware characteristics, thus maximizing performance transparently to applications. To achieve this goal, the data management community needs interdisciplinary collaboration with computer architecture, compiler and operating systems researchers. This involves rethinking traditional data structures, query processing algorithms, and database software architectures to adapt to the advances in the underlying hardware infrastructure.

4th International Workshop on Testing Database Systems
http://db.uwaterloo.ca/dbtest2011/

The DBTest workshops bring researchers and practitioners together to discuss key problems and ideas related to testing database systems. DBTest 2011 is the fourth workshop in the series. The focus of the DBTtest workshops is techniques for measuring important properties of database systems, including performance, reliability, security, availability, and robustness. We define database systems broadly to include any systems that, like relational database management systems, must manage a subtantial amount of data on behalf of applications.

NOTES SPACE

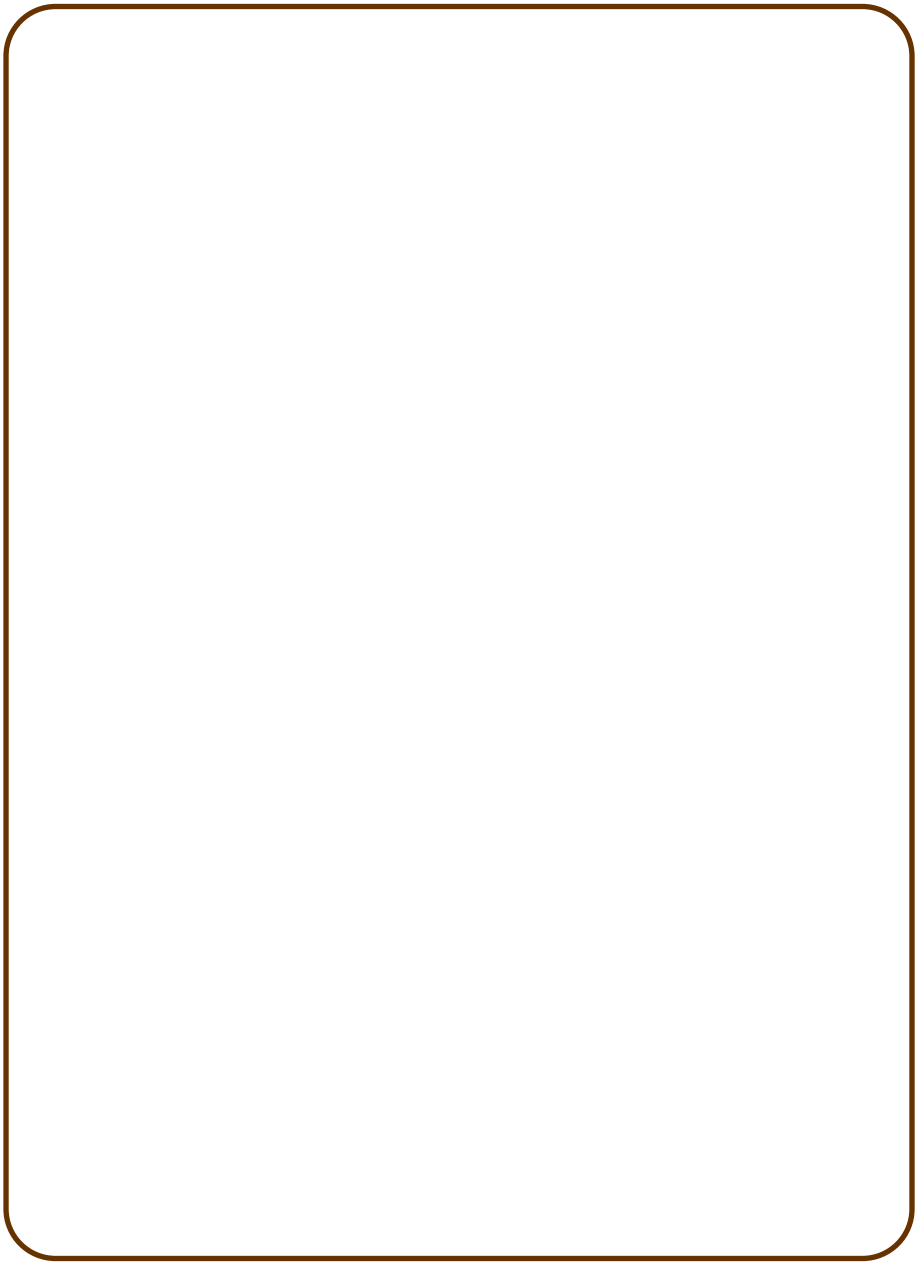

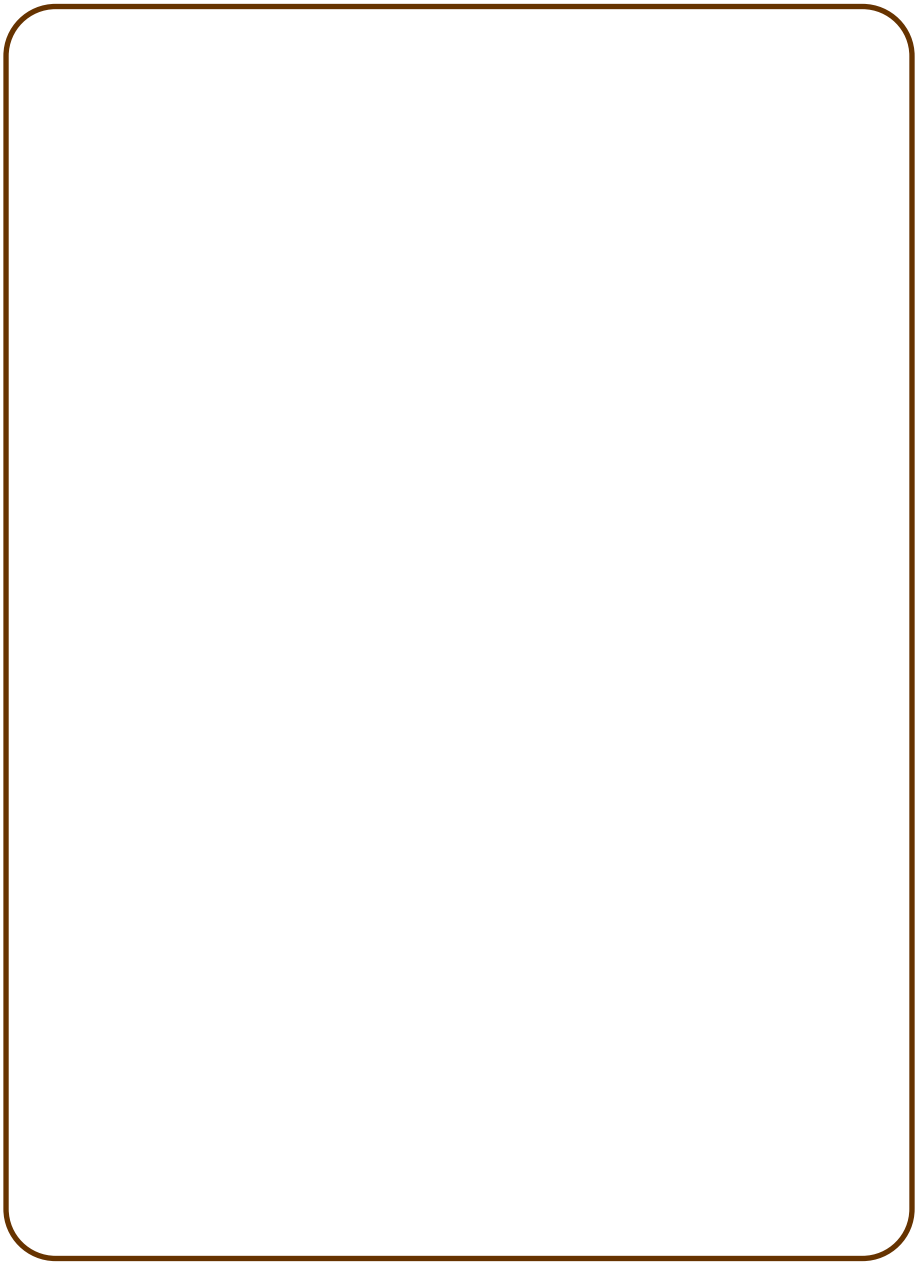

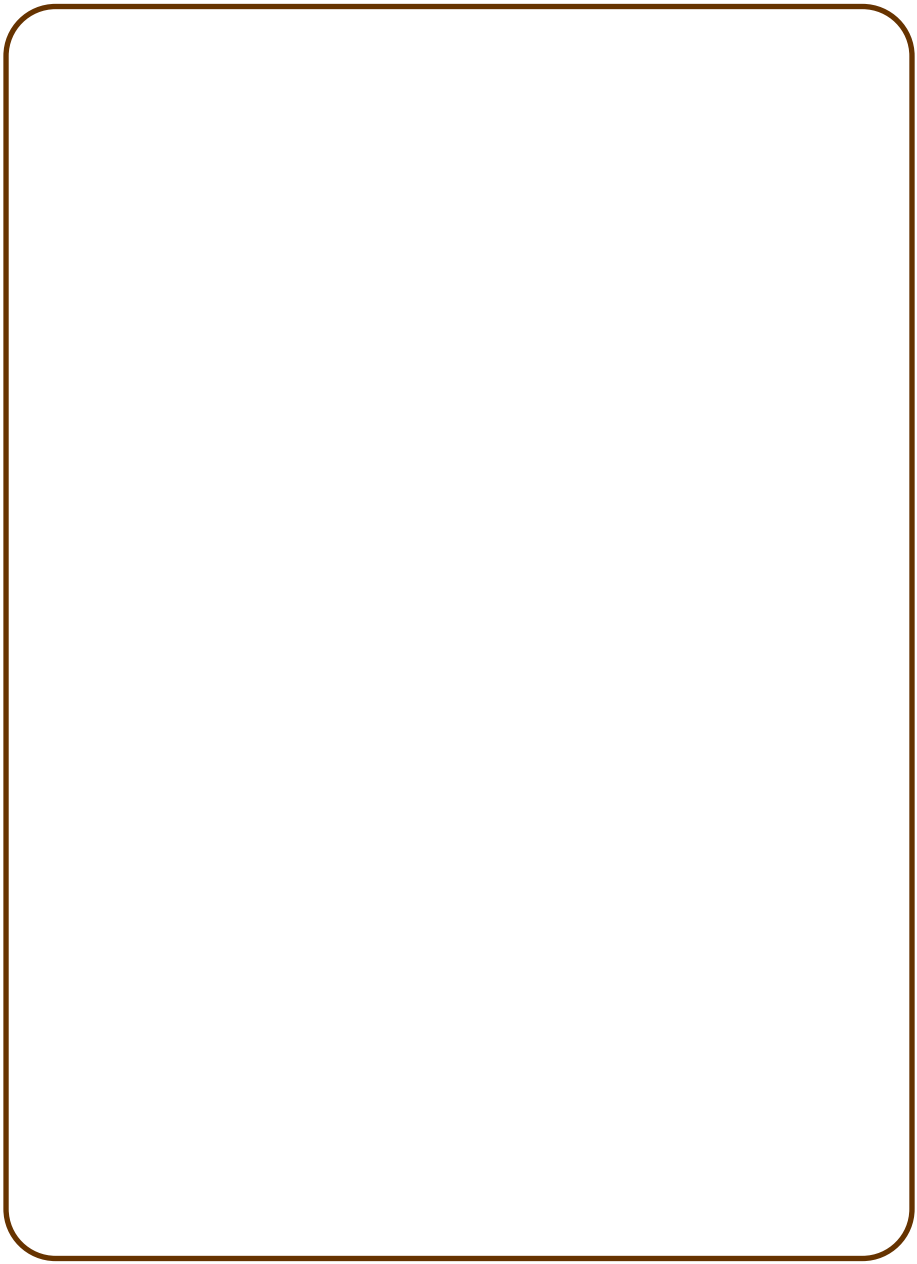

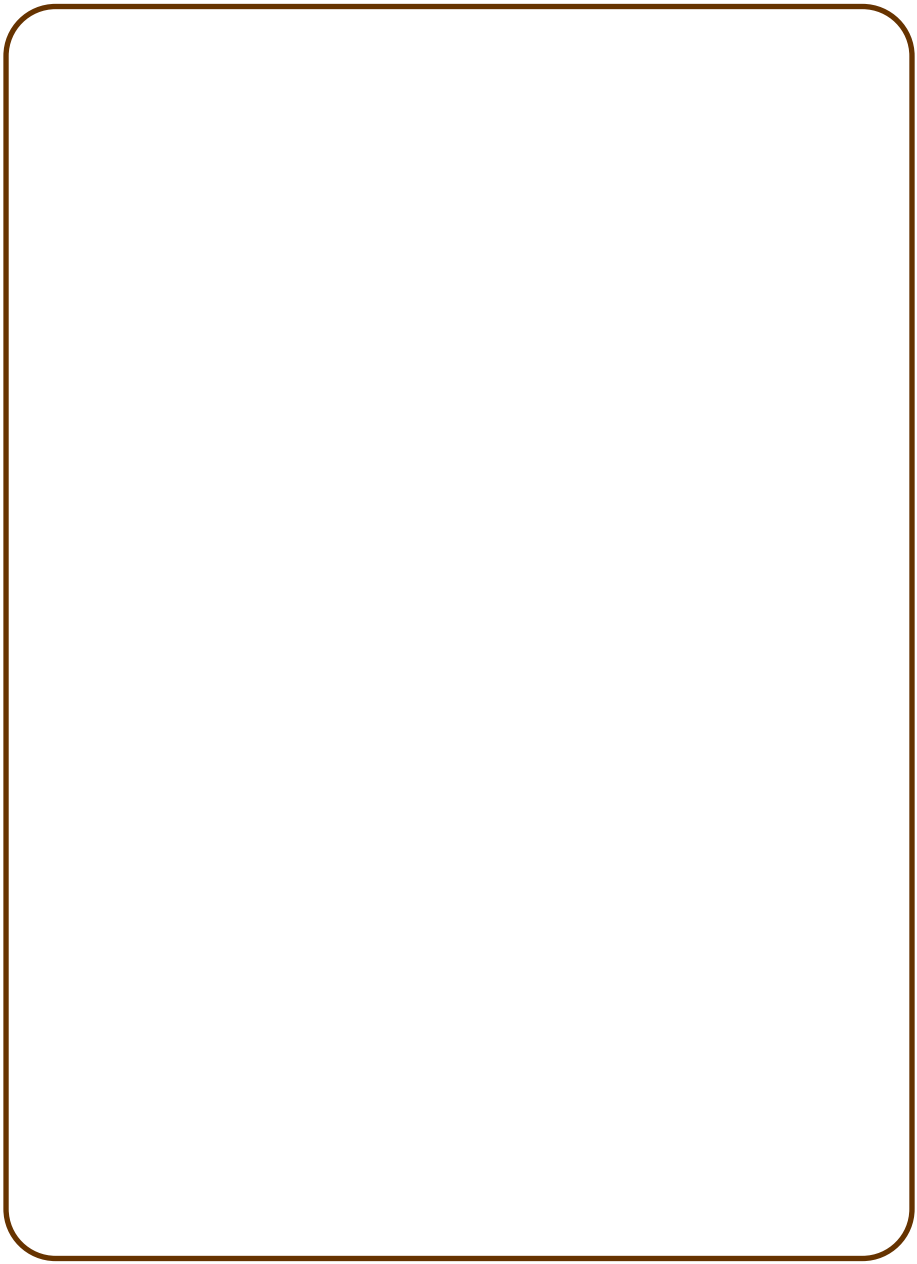

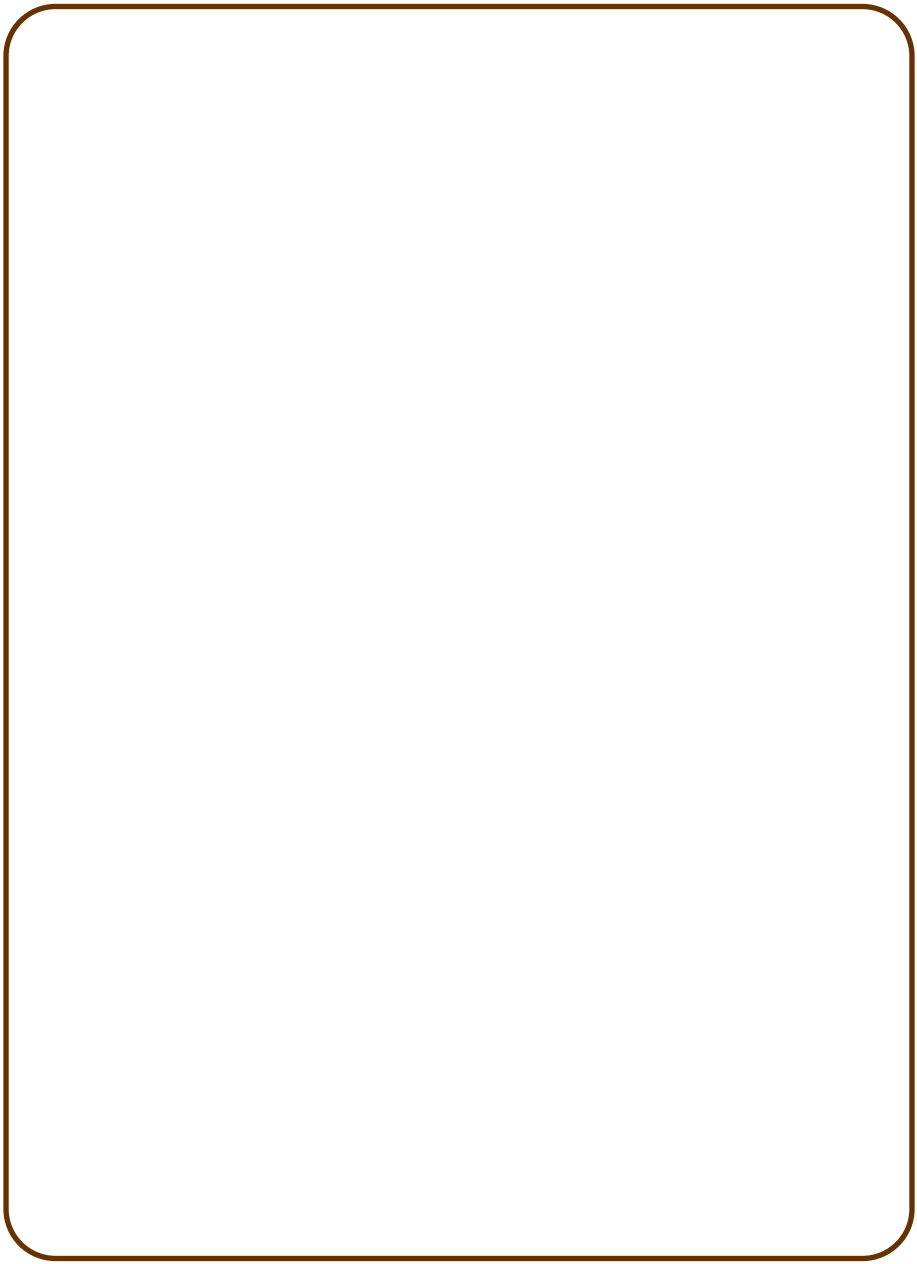

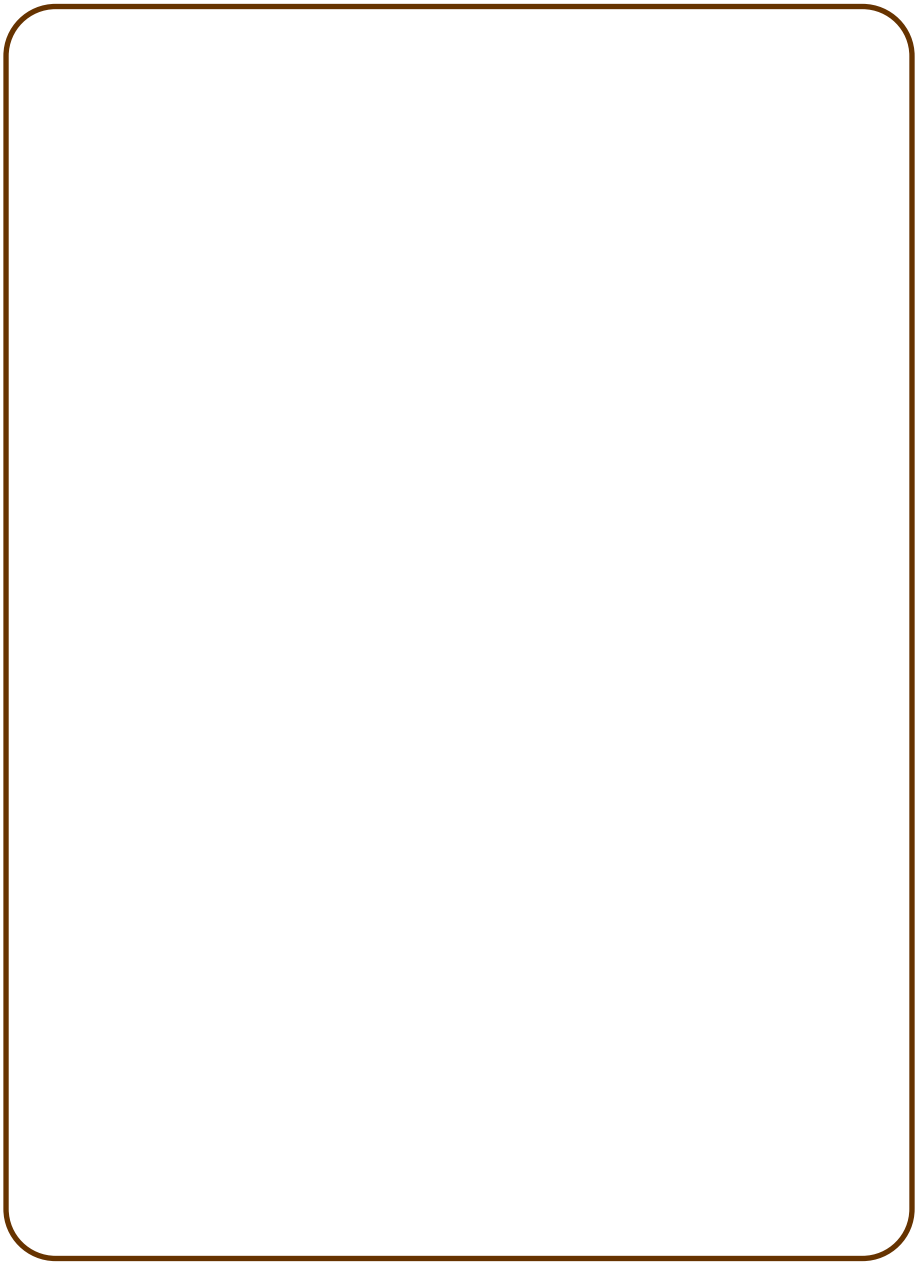

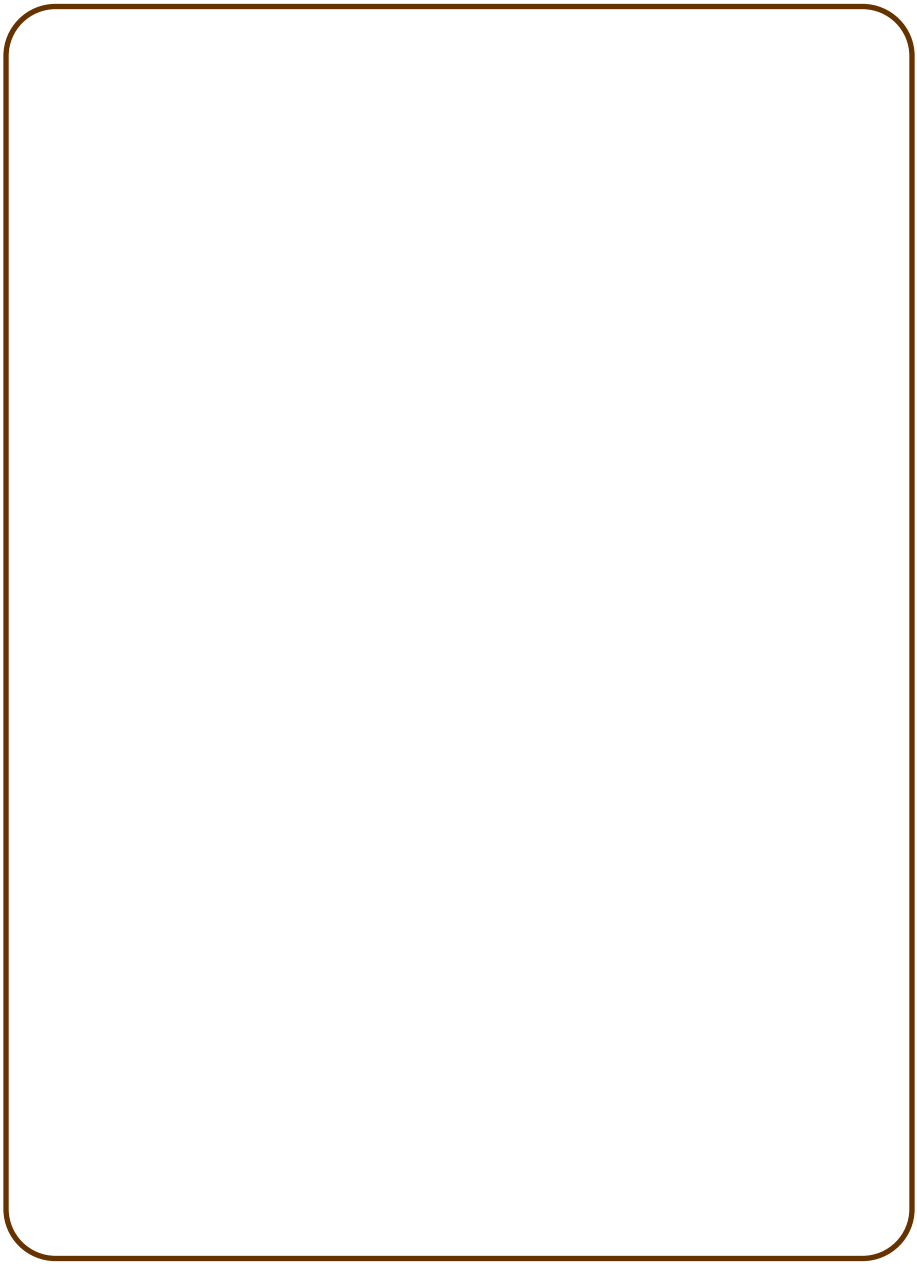

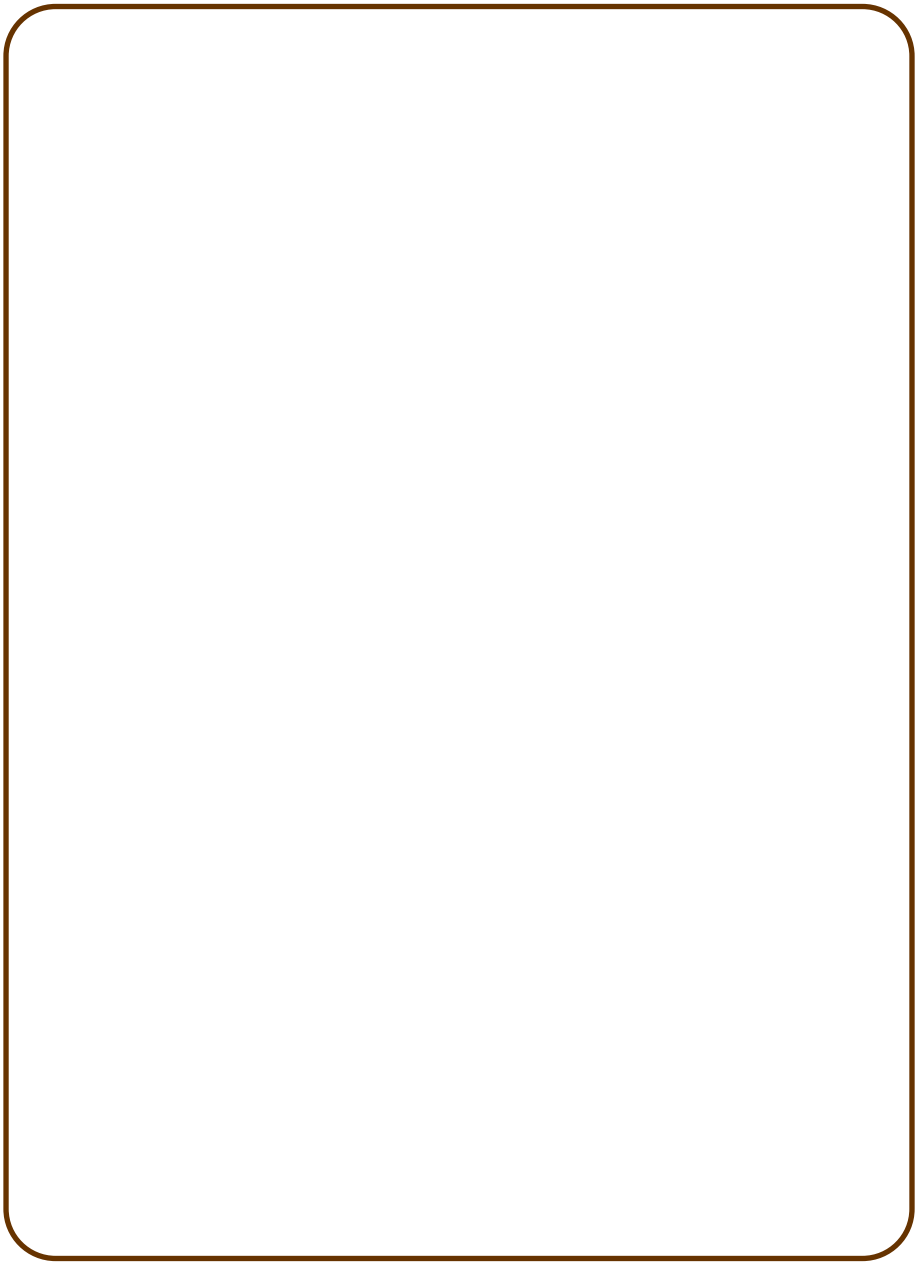

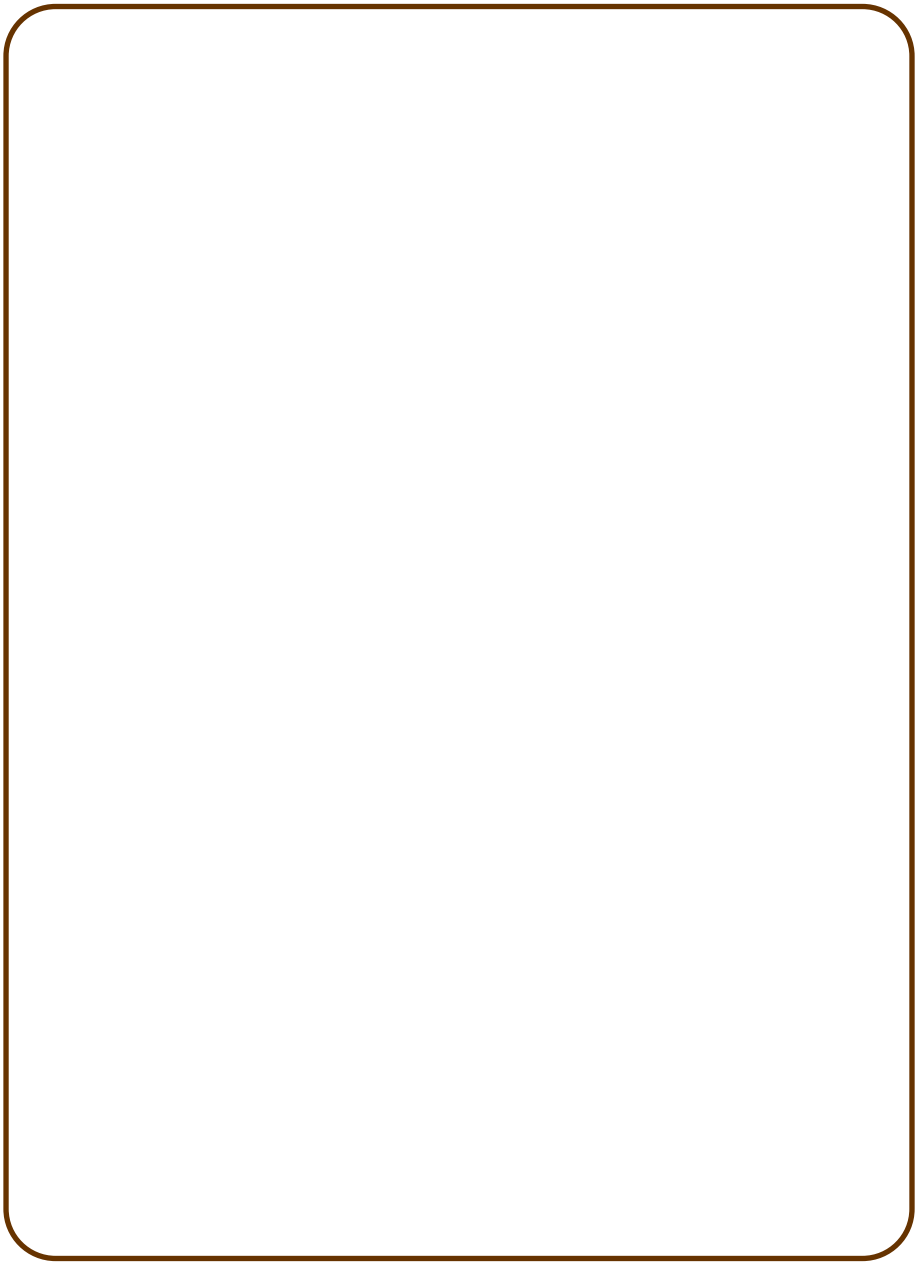

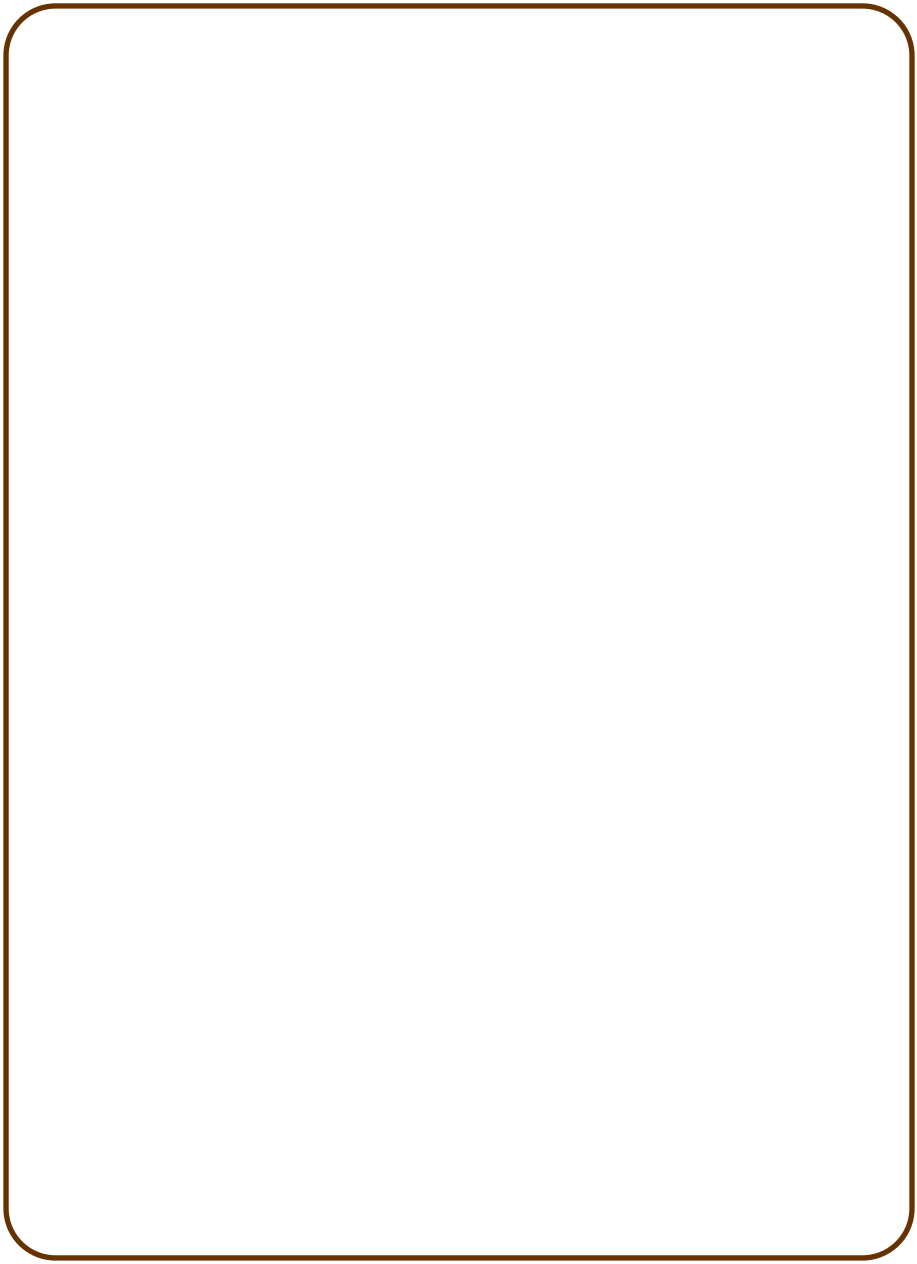

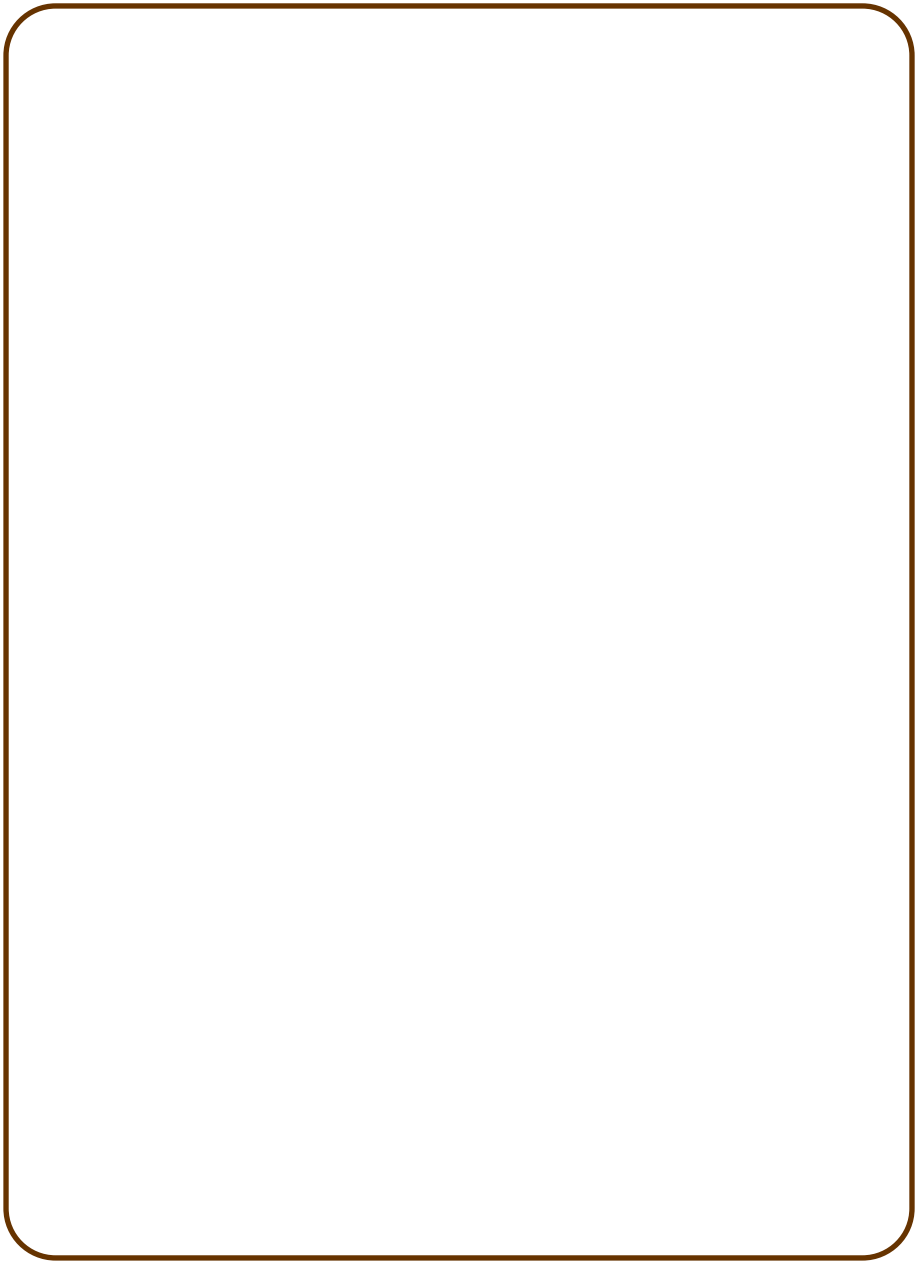

Sponsors:

Platinum Supporters:

Gold Supporters:

Silver Supporters:

Additional supporters: