ACM SIGMOD/PODS 2012

CONFERENCE PROGRAM

May 20-24, 2012
Scottsdale, Arizona, USA
http://www.sigmod.org/2012

ACM SIGMOD International Conference on Management of Data
&
ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems

SIGMOD/PODS 2012

May 20–24 2012
Scottsdale, Arizona, USA
http://www.sigmod.org/2012/

Conference Program

Editors: *K. Selçuk Candan (Arizona State University), Yi Chen (Arizona State University), Maurizio Lenzerini (University of Rome La Sapienza), Luis Gravano (Columbia University), Michael Benedikt (University of Oxford), Ariel Fuxman (Microsoft Research), Markus Krötzsch (University of Oxford)*

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

TABLE OF CONTENTS

Welcome Message from the SIGMOD Chairs

We are delighted to welcome you to SIGMOD 2012, the 2012 edition of the ACM SIGMOD International Conference on Management of Data, in Scottsdale, Arizona, in the Southwest of the United States. Scottsdale is in the heart of the Sonoran Desert and offers stunning desert vistas and a breathtaking setting for the conference. At the same time, Scottsdale is adjacent to Phoenix, one of the largest and fastest-growing cities in the United States.

SIGMOD 2012 hosts an exciting technical program, with two keynote talks, by Pat Hanrahan (Stanford University and Tableau Software) and Amin Vahdat (University of California, San Diego and Google); a plenary session with "Perspectives on Big Data," by Donald Kossmann (ETHZ), Kristen LeFevre (Google Research and University of Michigan), Sam Madden (MIT), and Anand Rajaraman (@WalmartLabs); 48 research paper presentations; six tutorials; 30 demonstrations; and 18 industrial presentations. In addition to having full 30-minute presentation slots, research papers are included in one of two Research Plenary Poster Sessions. One of these sessions is jointly for PODS and SIGMOD research papers, to deepen the ties between the two conferences. Another new plenary poster session, for papers from the 11 workshops co-located with SIGMOD 2012, is an effort to strengthen the link and synergy between the workshops and the conference.

SIGMOD 2012 includes several technical and social events designed specifically for student attendees. The SIGMOD/PODS 2012 Ph.D. Symposium, the Database Mentoring Workshop, the Undergraduate Research Poster Competition, and the New Researcher Symposium are all established components of the SIGMOD program and are all part of SIGMOD 2012. The conference also hosts a session dedicated to highlighting the finalists of the SIGMOD Programming Contest. (This year's task is to implement a multidimensional, high-throughput, in-memory indexing system.) In addition, the conference includes a new Information Session on Careers in Industry, aimed at bringing student attendees together with our Gold, Platinum, and Diamond sponsors, as well as "vis-à-vis" meetings aimed at helping Ph.D. students meet internationally recognized researchers in their research areas, to exchange ideas and receive guidance in a relaxed social setting.

We are immensely grateful to the many researchers who have shaped the conference program. In particular, we thank the Keynote and Panel Chair, Surajit Chaudhuri (Microsoft Research); the Tutorial Chair, Alon Halevy (Google Research); the Industrial Program Chair, AnHai Doan (University of Wisconsin-Madison and @WalmartLabs), together with the seven members of the Industrial Program Committee; the Demonstration Chair, Magdalena Balazinska (University of Washington), together with the 32 members of the Demonstration Program Committee; the Workshop Chair, Christian S. Jensen (Aarhus University); the Undergraduate Research Program Chair, Christopher Ré (University of Wisconsin-Madison); and the New Researcher Symposium Chairs, Xin (Luna) Dong (AT&T Labs–Research) and Christopher Ré (University of Wisconsin-Madison). We also express our enormous gratitude to the 88 members of the SIGMOD 2012 Program Committee, as well as to the 10 Program Committee "group leaders," Anastasia Ailamaki (EPFL), Philip Bernstein (Microsoft Research), Elisa Bertino (Purdue University), Umeshwar Dayal (HP Labs), Juliana Freire (NYU-Poly), Minos Garofalakis (Technical University of Crete), Donald Kossmann (ETHZ), Tova Milo (Tel Aviv University), Divesh Srivastava (AT&T Labs–Research), and Gerhard Weikum (Max-Planck Institute for Informatics). We received 289 research paper submissions; Program Committee members produced at least three reviews per research paper, and the group leaders expertly ensured that all papers received fair and thoughtful consideration through thorough reviews and discussion.

We also acknowledge the members of the several SIGMOD 2012 organizing committees, as well as the SIGMOD Executive Committee, for invaluable help and guidance throughout the many months leading to the conference. Lisa Singh deserves special thanks for providing us with her advice and SIGMOD institutional memory, as well as for her help with logistics in the early conference planning stages. We are also grateful to the Microsoft Research Conference Management Toolkit (CMT) team for their extremely prompt and helpful support throughout the complex conference reviewing process. (We used CMT for the submission and reviewing of research papers, demonstrations, and industrial presentations.) We also extend our appreciation to the student volunteers.

We thank the many SIGMOD 2012 supporters, whose contributions helped to maintain the conference registration fees low, most significantly for student attendees: Diamond Level: SAP; Platinum Level: Greenplum-EMC, IO Data Centers, Microsoft, and Oracle; Gold Level: Facebook, Google, IBM Research, LinkedIn, MarkLogic, Turn, Twitter, Vertica, and VMWare; and Silver Level: Actian, AT&T Labs–Research, HP, @WalmartLabs, and Yahoo! Arizona State University and the University of Arizona have also extended their generous support to SIGMOD 2012. Finally, the National Science Foundation and the SIGMOD Executive Committee have financially supported the SIGMOD 2012 student travel award program.

We welcome you to Scottsdale and we sincerely hope that you will enjoy SIGMOD 2012!

K. Selçuk Candan	**Yi Chen**	**Luis Gravano**	**Ariel Fuxman**
General Chair	*General Chair*	*Program Chair*	*Proceedings Chair*

Welcome Message from the PODS General Chair

It is our great pleasure to welcome you to the 2012 ACM Symposium on Principles of Database Systems – PODS'12, held in Scottsdale, Arizona, USA on May 20–23, 2012, in conjunction with the 2012 ACM SIGMOD International Conference on Management of Data.

This year's symposium continues its tradition of being the premier international conference on the theoretical aspects of data management. Since the first edition of the symposium in 1982, the 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 interested reader is referred to the PODS web pages at http://www09.sigmod.org/sigmod/pods/ for various information on the history of the conference series.

As usual, putting together PODS'12 was a team effort. We are particularly grateful to the Program Chair, Michael Benedikt, who did a magnificent job in selecting and coordinating the program committee members, and to the whole program committee, who worked very hard in reviewing papers and providing feedback for authors. We are also grateful to the SIGMOD 2012 General Chairs, K. Selçuk Candan and Yi Chen, for their collaboration in all the issues requiring coordination between SIGMOD and PODS. Finally, we thank Markus Krötzsch, the PODS'12 Proceedings and Publicity Chair, Wim Martens, for maintaining the PODS web pages, and all our sponsors, in particular the ACM Special Interest Groups on Management of Data, for their invaluable support.

We wish you a profitable and enjoyable stay in Arizona, and we hope that you will find the PODS'12 program exciting and thought provoking, in the best tradition of the PODS Symposium.

Maurizio Lenzerini
PODS'12 General Chair

Welcome Message from the PODS Program Chair

First, a brief overview of the contents of this volume, the proceedings of the thirty-first ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2012). The proceedings includes a paper based on the keynote address by Surajit Chaudhuri along with two papers based on invited tutorials, one by Michael Mahoney and one by Benjamin Pierce. There are 26 research papers that were selected by the Program Committee, out of 101 submissions with authors from over 25 countries across the world. Out of the 26 accepted papers, the program committee selected the paper Worst Case Optimal Join Algorithms by Hung Q. Ngo, Ely Porat, Christopher Ré and Atri Rudra for the PODS 2012 Best Paper Award. In addition, the announcement of the 2012 ACM PODS Alberto O. Mendelzon Test-of-Time Award appears in the proceedings, given to Containment and Equivalence for an XPath Fragment by Gerome Miklau and Dan Suciu. The latter paper originally appeared in the proceedings of PODS 2002. Congratulations to the authors of these papers.

The review process was grueling, and involved enormous effort from a large group of researchers. This year PODS experimented with the use of an External Review Committee, consisting of distinguished experts in areas of particular interest to PODS, in addition to our core Program Committee. We relied heavily on the EasyChair system for management of all aspects of the review process, and we are extremely grateful to Andrei Voronkov for his help in adding and modifying new features to support the External Review Committee. All members of the Program Committee, External Review Committee, and the additional external referees deserve thanks for their work – both for producing the selection of papers that appear here in the proceedings, and for providing high-quality feedback to all authors of submitted papers.

The proceedings would not have been possible without the constant attention and support of Maurizio Lenzerini, the PODS General Chair. The PODS Executive Committee helped select the Program Committee and External Review Committee, and played an important role in advising on issues of policy during the conference. Special thanks are also due to Markus Krötzsch, who served as the PODS Proceedings Chair and as PODS Publicity Chair – both of these jobs were more complex than usual this year due to the earlier date of the conference. In advising us we leaned heavily on the advice of Thomas Schwentick and Wim Martens, the previous PODS Program Chair and Publicity Chair, who patiently answered questions on every aspect of the conference.

Many people outside of the PODS community also provided critical support. I want to particularly express heartfelt thanks to K. Selçuk Candan and Yi Chen, the SIGMOD General Chairs: great colleagues who supported PODS in every possible way. We also thank Huiping Cao and Yan Qi, the Web/Information Chairs, for their great work in managing the website.

Finally, I thank the SIGMOD Executive Committee for their help in arranging supplemental support from EasyChair for the 2012 electronic review process, and the SIGMOD/PODS sponsors for their support.

Michael Benedikt
PODS'12 Program Chair

SIGMOD CONFERENCE ORGANIZATION

General Chairs:	K. Selçuk Candan (Arizona State University) Yi Chen (Arizona State University)
General Vice-Chair:	Richard Snodgrass (University of Arizona)
Program Chair:	Luis Gravano (Columbia University)
Proceedings Chair:	Ariel Fuxman (Microsoft Research)
Tutorial Chair:	Alon Halevy (Google Research)
Keynote and Panel Chair:	Surajit Chaudhuri (Microsoft Research)
Industrial Program Chair:	AnHai Doan (University of Wisconsin-Madison and @WalmartLabs)
Demonstration Chair:	Magdalena Balazinska (University of Washington)
Workshop Chair:	Christian S. Jensen (Aarhus University)
Undergraduate Research Program Chair:	Christopher Ré (University of Wisconsin-Madison)
New Researcher Symposium Chairs:	Xin (Luna) Dong (AT&T Labs-Research) Christopher Ré (University of Wisconsin-Madison)
Finance Chairs:	Egemen Tanin (University of Melbourne) Junichi Tatemura (NEC Labs)
Publicity/Social Media Chairs:	Lei Chen (HKUST) Maria Luisa Sapino (University of Torino)
Sponsorship Chairs:	Divyakant Agrawal (University of California at Santa Barbara) Vasilis Vassalos (Athens U. of Economics and Business)
Exhibit Chairs:	Wen-Syan Li (SAP) Berthold Reinwald (IBM Almaden Research Center)

Local Arrangements Chair:	Suzanne Dietrich (Arizona State University)
Registration Chairs:	Ziyang Liu (NEC Labs) Bongki Moon (University of Arizona)
Demonstration and Workshop Local Arrangements Chair:	Hasan Davulcu (Arizona State University)
Mentorship Chairs:	Feifei Li (University of Utah) Nan Zhang (George Washington University)
Web/Information Chairs:	Huiping Cao (New Mexico State University) Yan Qi (Turn Inc.)
Program Committee Group Leaders:	Anastasia Ailamaki (EPFL) Philip Bernstein (Microsoft Research) Elisa Bertino (Purdue University) Umeshwar Dayal (HP Labs) Juliana Freire (NYU-Poly) Minos Garofalakis (Technical University of Crete) Donald Kossmann (ETHZ) Tova Milo (Tel Aviv University) Divesh Srivastava (AT&T Labs-Research) Gerhard Weikum (Max-Planck Institute for Informatics)
Program Committee:	Ashraf Aboulnaga (University of Waterloo) Arvind Arasu (Microsoft Research) Marcelo Arenas (PUC Chile) Shivnath Babu (Duke University) Pablo Barceló (Universidad de Chile) Philip Bohannon (Yahoo! Research) Angela Bonifati (CNR) Michael Cafarella (University of Michigan, Ann Arbor) Malu Castellanos (Hewlett-Packard Laboratories) Kevin Chen-Chuan Chang (UIUC) Yi Chen (Arizona State University) Rada Chirkova (North Carolina State University) Junghoo Cho (University of California, Los Angeles) Christopher Clifton (Purdue University) Nilesh Dalvi (Yahoo! Research) Anish Das Sarma (Google Research) Atish Das Sarma (Google Research) Amol Deshpande (University of Maryland) Yanlei Diao (University of Massachusetts, Amherst) Xin Luna Dong (AT&T Labs-Research) Alan Fekete (University of Sydney) Helena Galhardas (Technical University of Lisbon) Venkatesh Ganti (Google) Phillip Gibbons (Intel Labs) Aristides Gionis (Yahoo! Research)

Program Committee (continued):

Ashish Goel (Stanford University)
Sudipto Guha (University of Pennsylvania)
Zoltán Gyöngyi (Google Research)
Marios Hadjieleftheriou (AT&T Labs-Research)
Wook-Shin Han (Kyungpook National University)
Vagelis Hristidis (University of California, Riverside)
Stratos Idreos (CWI)
Zachary Ives (University of Pennsylvania)
H. V. Jagadish (University of Michigan, Ann Arbor)
Christopher Jermaine (Rice University)
Ryan Johnson (University of Toronto)
Theodore Johnson (AT&T Labs-Research)
Panos Kalnis (KAUST)
Benjamin Kao (The University of Hong Kong)
Alfons Kemper (TU München)
Martin Kersten (CWI)
Christoph Koch (EPFL)
George Kollios (Boston University)
Henry Korth (Lehigh University)
Alexandros Labrinidis (University of Pittsburgh)
Kristen LeFevre (University of Michigan, Ann Arbor)
Jure Leskovec (Stanford University)
Guy Lohman (IBM Almaden Research Center)
David Lomet (Microsoft Research)
Qiong Joan Luo (HKUST)
Ashwin Machanavajjhala (Yahoo! Research)
Ioana Manolescu (INRIA)
Volker Markl (TU Berlin)
Gerome Miklau (University of Massachusetts, Amherst)
Mirella Moro (Universidade Federal de Minas Gerais)
Vivek Narasayya (Microsoft Research)
Frank Olken (Lawrence Berkeley National Laboratory)
Beng Chin Ooi (National University of Singapore)
Fatma Özcan (IBM Almaden Research Center)
Olga Papaemmanouil (Brandeis University)
Yannis Papakonstantinou (University of California, San Diego)
Stelios Paparizos (Microsoft Research)
Jorge Pérez (Universidad de Chile)
Evaggelia Pitoura (University of Ioannina)
Cecilia Procopiuc (AT&T Labs-Research)
Erhard Rahm (University of Leipzig)
Jun Rao (LinkedIn)
Rajeev Rastogi (Yahoo! Labs)
Kenneth Ross (Columbia University)
Michael Rys (Microsoft)
Cyrus Shahabi (University of Southern California)
Kyuseok Shim (Seoul National University)
Narayanan Shivakumar (1618Labs, Inc.)
Lisa Singh (Georgetown University)
Radu Sion (Stony Brook University)
Dan Suciu (University of Washington)
S. Sudarshan (IIT Bombay)
Wang-Chiew Tan (University of California, Santa Cruz)
Val Tannen (University of Pennsylvania)
Yufei Tao (Chinese University of Hong Kong)
Martin Theobald (Max-Planck Institute for Informatics)
Yuanyuan Tian (IBM Almaden Research Center)
Vasilis Vassalos (Athens University of Economics and Business)
Stratis Viglas (University of Edinburgh)

Program Committee (continued):

Wei Wang (University of North Carolina at Chapel Hill)
Raymond Chi-Wing Wong (HKUST)
Jun Yang (Duke University)
Cong Yu (Google Research)

Demonstration Program Committee:

Yanif Ahmad (Johns Hopkins University)
Vinayak Borkar (University of California, Irvine)
Lei Chen (Hong Kong University of Science and Technology)
Hong Cheng (The Chinese University of Hong Kong)
Tyson Condie (Yahoo! Research)
Ada Fu (Chinese University of Hong Kong)
Wolfgang Gatterbauer (Carnegie Mellon University)
Shenoda Guirguis (Intel)
Michael Hay (Cornell University)
Thomas Heinis (EPFL)
Bill Howe (University of Washington)
Jeong-Hyon Hwang (State University of New York at Albany)
Yannis Kotidis (Athens University of Economics and Business)
Wang-Chien Lee (Pennsylvania State University)
Eric Lo (The Hong Kong Polytechnic University)
Samuel Madden (MIT)
David Maier (Portland State University)
Alexandra Meliou (University of Washington)
Mohamed Mokbel (University of Minnesota)
Carlos Ordonez (University of Houston)
Thanasis Papaioannou (EPFL)
Rachel Pottinger (University of British Columbia)
Ravi Ramamurthy (Microsoft Research)
Vibhor Rastogi (Yahoo! Research)
Mirek Riedewald (Northeastern University)
Russel Sears (Yahoo! Research)
Nesime Tatbul (ETH Zurich)
Boon Thau Loo (University of Pennsylvania)
Marcos Vaz Salles (University of Copenhagen)
Daisy Zhe Wang (University of Florida)
Demetris Zeinalipour (University of Cyprus)
Jingren Zhou (Microsoft)

Industrial Program Committee:

David DeWitt (Microsoft)
Jayant Madhavan (Google)
Ullas Nambiar (IBM Research)
Chris Olston (Bionica Human Computing)
Anand Rajaraman (@WalmartLabs)
Raghu Ramakrishnan (Yahoo! Research)
Shivakumar Vaithyanathan (IBM Almaden Research Center)

Travel Award Selection Committee:

Graham Cormode (AT&T Labs-Research)
Tingjian Ge (University of Massachusetts, Lowell)
Murat Kantarcioglu (University of Texas at Dallas)
Anastasios Kementsietsidis (IBM Research)
Alexandros Labrinidis (University of Pittsburgh)
Chengkai Li (University of Texas at Arlington)
Feifei Li (University of Utah)
Ashwin Machanavajjhala (Yahoo! Research)
Senjuti Basu Roy (University of Washington Tacoma)
Nan Zhang (George Washington University)

PODS CONFERENCE ORGANIZATION

General Chair: Maurizio Lenzerini (University of Rome La Sapienza)

Program Chair: Michael Benedikt (University of Oxford)

Proceedings & Publicity Chair: Markus Krötzsch (University of Oxford)

Program Committee:
Mikhail Atallah (Purdue University)
Toon Calders (Technische Universiteit Eindhoven)
Diego Calvanese (Free University of Bozen-Bolzano)
James Cheney (University of Edinburgh)
Graham Cormode (AT&T Labs Research)
Alin Deutsch (University of California, San Diego)
Gianluigi Greco (Universita della Calabria)
Todd J. Green (University of California, Davis)
Martin Grohe (Humboldt-Universität zu Berlin)
Marc Gyssens (Hasselt University)
T.S. Jayram (IBM Research – Almaden & IBM Research – India)
Daniel Kifer (Penn State University)
Phokion Kolaitis (University of California, Santa Cruz & IBM Research – Almaden)
Rasmus Pagh (IT University of Copenhagen)
Luc Segoufin (INRIA Cachan)
Pierre Senellart (Télécom ParisTech)
Sophie Tison (Lille University)
Victor Vianu (University of California, San Diego)
David Woodruff (IBM Research – Almaden)

External Review Committee:

Querying and Mining of Unstructured Data

AnHai Doan (University of Wisconsin)
Aristides Gionis (Yahoo! Research Barcelona)
Djoerd Hiemstra (University of Twente)
Stefano Leonardi (University of Rome La Sapienza)
Evimaria Terzi (Boston University)

Web Services, Web Programming and Data-Centric Workflow

Wil van der Aalst (Technische Universiteit Eindhoven)
Anders Møller (Aarhus University)
Farouk Toumani (ISIMA & Blaise Pascal University)
David Walker (Princeton University)
Karsten Wolf (University of Rostock)

Learning of Data Models and Queries

Deepak Agarwal (Yahoo! Research Silicon Valley)
James Cussens (York University)
Amol Deshpande (University of Maryland)
Kristian Kersting (Fraunhofer Institute IAIS & University of Bonn)

External Review Committee (continued):

Cloud Computing and Next-generation Distributed Query Processing

Shivnath Babu (Duke University)
Phillip Gibbons (Intel Labs Pittsburgh)
Monica Lam (*Stanford* University)
Boon Thau Loo (University of Pennsylvania)
Volker Markl (Technische Universität Berlin)

Semantic, Linked, Networked, and Crowdsourced Data

Panagiotis Ipeirotis (New York University)
David Karger (Massachusetts Institute of Technology)
Carsten Lutz (University of Bremen)
Boris Motik (University of Oxford)

Additional Reviewers:

Peyman Afshani
Alexandr Andoni
Marcelo Arenas
Pablo Barceló
Ioan Marius Bilasco
Peter Boncz
Jan Van den Bussche
Hubert T.H. Chan
Anne-Cécile Caron
Ken Clarkson
Reynold Cheng
Amit Chakrabarti
Venkatesan Chakaravarthy
Sara Cohen
Prasad Deshpande
Anindya De
Thierry Delot
Daniel Dougherty
Cynthia Dwork
Wolfgang Faber
Sergio Flesca
Michael E. Houle
Sumit Ganguly
Vikas Garg
Bart Goethals
Dimitrios Gunopulos
Roberto Grossi
Moritz Hardt
Sariel Har-Peled
Fritz Henglein
André Hernich
Lam Thanh Hoang
Wing-Kai Hon
Satyen Kale
Murat Kantarcioglu
Grigoris Karvounarakis
Shiva Kasiviswanathan
Jonathan Kelner
Sreyash Kenkre
Evgeny Kharlamov
Benny Kimelfeld
Bart Kuijpers
Janardhan Kulkarni
Daniel Lemire
Kasper Green Larsen
Dominique Laurent
Jan Vondrak
Jef Wijsen
Ryan Williams
Ranko Lazic
Domenico Lembo
Maurizio Lenzerini
Bertram Ludaescher
Edward Lui
S. Muthukrishnan
Sebastian Maneth
Dániel Marx
Anca Muscholl
Andrew McGregor
Ashwin Machanavajjhala
Gerome Miklau
Ramasuri Narayanam
Gonzalo Navarro
Frank Neven
Joachim Niehren
Kobbi Nissim
Dan Olteanu
Martin Otto
Vinayaka Pandit
Jean-Marc Petit
Jeff Phillips
Pawel Parys
Debmalya Panigrahi
Mihai Patrascu
Vibhor Rastogi
Romain Rouvoy
Paolo Romano
Riccardo Rosati
Sambuddha Roy
Srinivasa Rao Satti
Mukund Sundararajan
Yogish Sabharwal
Nicole Schweikardt
Slawek Staworko
Francesco Scarcello
Dan Suciu
Yufei Tao
Giorgio Terracina
Srikanta Tirthapura
Marc Tommasi
Jonathan Ullman
Gregory Valiant
Philippe Vanheeghe
Stijn Vansummeren
Sergei Vassilvitzkii
Ke Yi
Hongbo Zhang
Qin Zhang

Alberto O. Mendelzon Test-of-Time Award Committee:

Richard Hull (*chair*, IBM T. J. Watson Research Center)
Phokion Kolaitis (University of California, Santa Cruz & IBM Research – Almaden)
Dirk Van Gucht (Computer Science Department, Indiana University)

SPONSORS, SUPPORTERS, AND EXHIBITORS

ACM Sponsor:

Additional ACM Sponsors for PODS:

Diamond:

Platinum:

Gold:

Silver:

Academic:

Student Travel Awards:

Undergraduate Competition and Jim Gray Dissertation Awards:

Student Support:

Exhibitor:

Platinum Publisher:

Gold Publisher:

Silver Publisher:

CONFERENCE VENUE

Hyatt Regency Scottsdale Resort and Spa at Gainey Ranch

The location of the 2012 ACM SIGMOD/PODS Conference is the Hyatt Regency Scottsdale Resort and Spa at Gainey Ranch, in Scottsdale, Arizona, USA. The amenities of this 27-acre hotel resort include:

- A 2.5-acre "water playground" consisting of 10 swimming pools, a three-story water slide, a sand beach, a whirlpool spa, and two poolside bars
- Spa Avania, with an outdoor heated mineral pool, state-of-the-art exercise facilities, steam and sauna
- Jogging paths and bicycling trails with onsite bicycles
- Golf courses at the 27-hole Gainey Ranch Golf Club, adjacent to the resort
- Four lighted tennis courts.

DIRECTIONS

From Phoenix Sky Harbor Int'l Airport: Exit the airport following the signs to 202 East. Take the 202 East to 101 North. Continue on 101 North to the Via De Ventura Exit and make a left. Travel West for approximately 2 1/2 miles (Via De Ventura turns into Doubletree Ranch Road). The Hyatt Regency Scottsdale Resort and Spa will be on the right hand side just before Scottsdale Road.

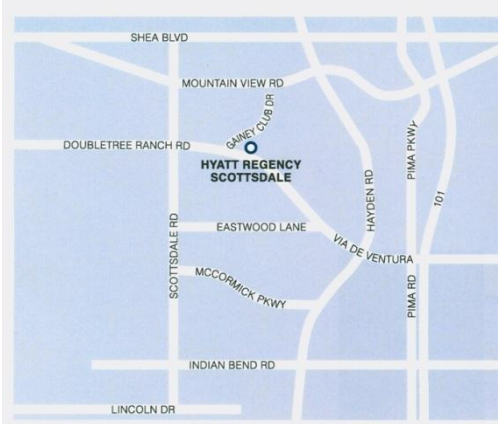

FLOOR PLAN

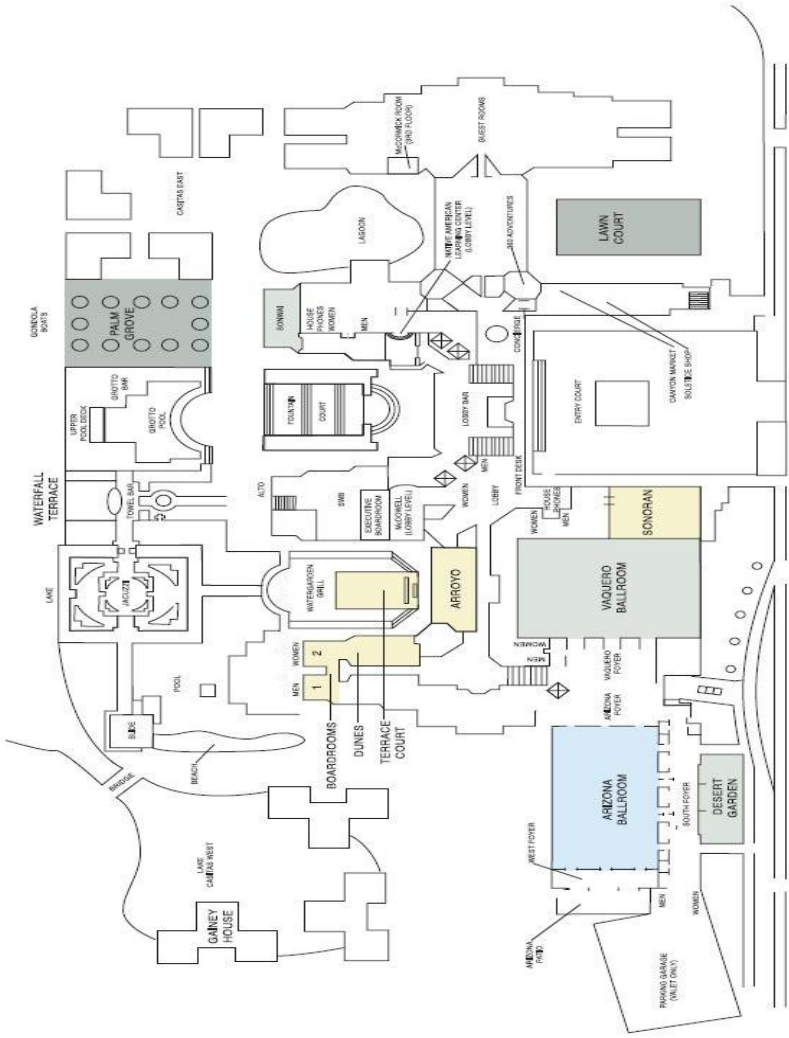
PALM GROVE
GROTTO BAR
GROTTO POOL
UPPER POOL DECK
WATERFALL TERRACE
TOWEL BAR
JACUZZI
POOL
BAR
BEACH
BRIDGE
LAKE
GONDOLA BOATS
CASITAS EAST
LAKE CASITAS WEST
GAINEY HOUSE
BOARDROOMS
DUNES
TERRACE COURT
WATERGARDEN GRILL
ARROYO
SONORAN
VAQUERO BALLROOM
VAQUERO FOYER
ARIZONA FOYER
ARIZONA BALLROOM
WEST FOYER
SOUTH FOYER
DESERT GARDEN
LAWN COURT
ENTRY COURT
CANYON MARKET / SOUVENIR SHOP
FRONT DESK
LOBBY
LOBBY BAR
CONCIERGE
FOUNTAIN COURT
LAGOON
SONORAN
HOUSE PHONES
EXECUTIVE BOARDROOM
McDOWELL (LOBBY LEVEL)
MACDONALD ROOM (3RD FLOOR)
GUEST ROOMS
NATIVE AMERICAN LEARNING CENTER (LOBBY LEVEL)
KID ADVENTURES
PARKING GARAGE (VALET ONLY)
MEN
WOMEN

FIRST FLOOR

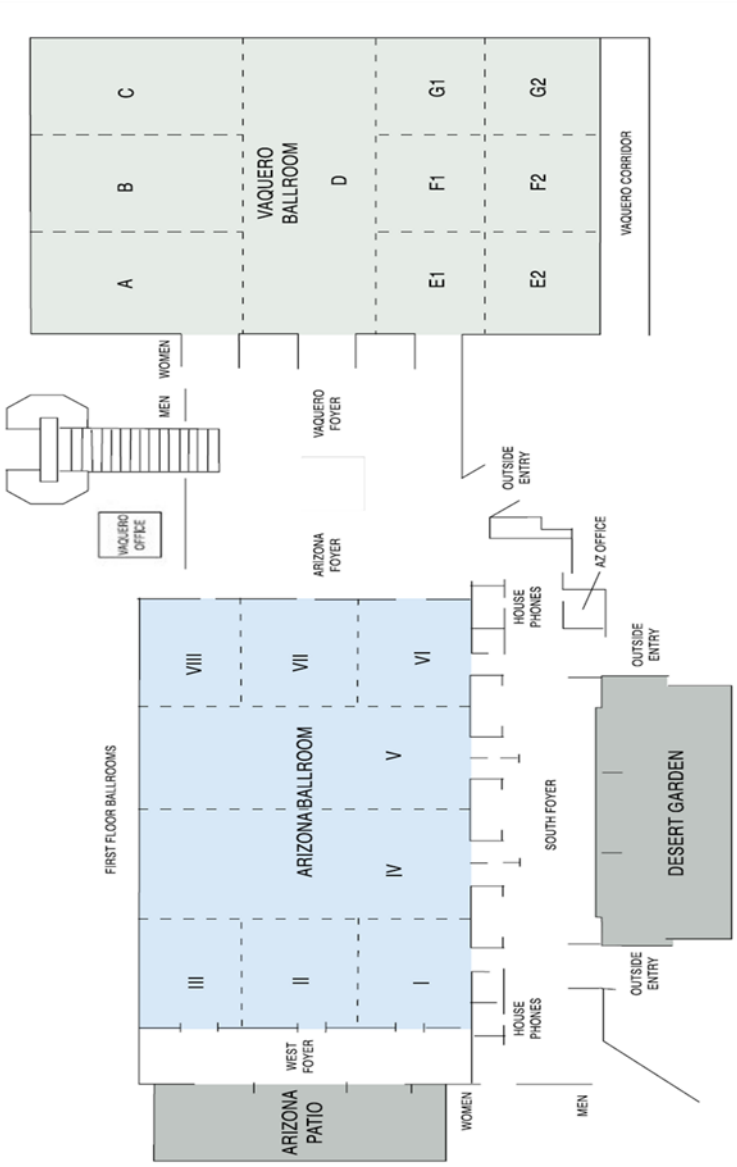
A
B
C
VAQUERO BALLROOM
D
E1
F1
G1
E2
F2
G2
VAQUERO CORRIDOR
WOMEN
MEN
VAQUERO FOYER
VAQUERO OFFICE
ARIZONA FOYER
OUTSIDE ENTRY
AZ OFFICE
HOUSE PHONES
OUTSIDE ENTRY
VIII
VII
VI
V
IV
III
II
I
FIRST FLOOR BALLROOMS
ARIZONA BALLROOM
SOUTH FOYER
DESERT GARDEN
OUTSIDE ENTRY
HOUSE PHONES
WEST FOYER
ARIZONA PATIO
WOMEN
MEN

SECOND FLOOR / ENTRY LEVEL

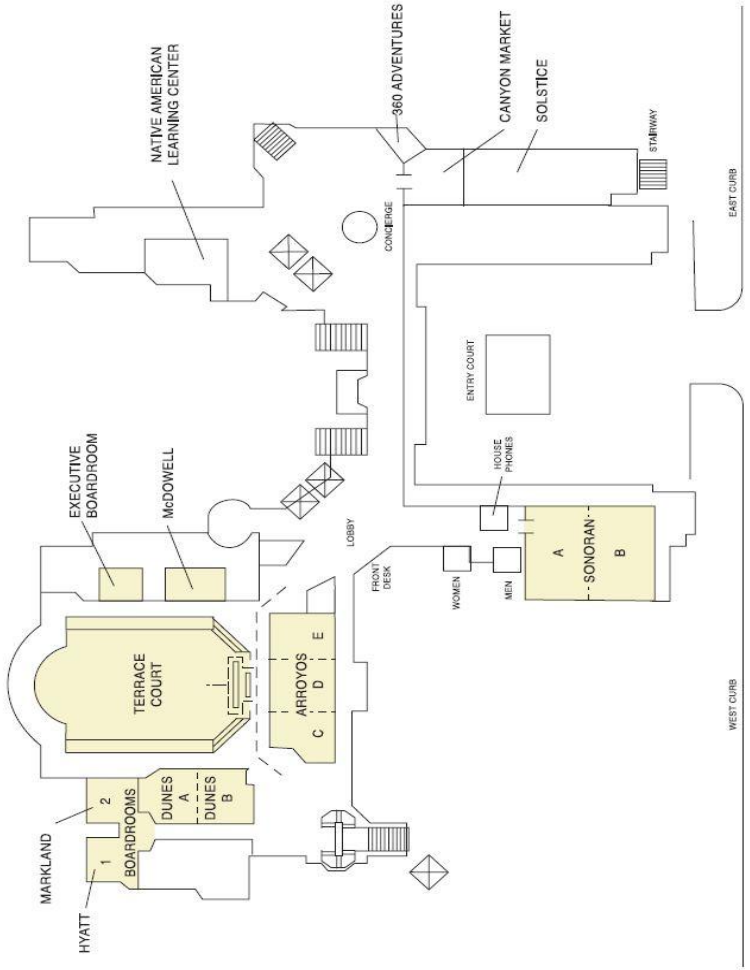

SECOND FLOOR / ENTRY LEVEL

PROGRAM AT A GLANCE

SUNDAY									
08:00- 08:30	**Continental Breakfast** Arizona, Vaquero, and South Foyers								
08:30- 10:00	**Ph.D. Symp.** Dunes A-B	**DB Me** Arizona VIII	**DBSocial** Arizona II	**IIWeb** Arizona VII	**KEYS** Arizona VI	**MobiDE** Arizona I	**SWEET** Arizona V	**SWIM** Arizona III	**WebDB** Arizona IV
10:00- 10:30	**Coffee Break** Arizona, South and West Foyers								
10:30- 12:00	**Ph.D. Symp.** Dunes A-B	**DB Me** Arizona VIII	**DBSocial** Arizona II	**IIWeb** Arizona VII	**KEYS** Arizona VI	**MobiDE** Arizona I	**SWEET** Arizona V	**SWIM** Arizona III	**WebDB** Arizona IV
12:00- 13:30	**Lunch Provided for Sunday Workshop Participants** Vaquero E-F								
13:30- 15:00	**Ph.D. Symp.** Dunes A-B	**DB Me** Arizona VIII	**DBSocial** Arizona II	**IIWeb** Arizona VII	**KEYS** Arizona VI	**MobiDE** Arizona I	**SWEET** Arizona V	**SWIM** Arizona III	**WebDB** Arizona IV
15:00- 15:30	**Coffee Break** Arizona, South and West Foyers								
15:30- 17:00	**Ph.D. Symp.** Dunes A-B	**DB Me** Arizona VIII	**DBSocial** Arizona II	**IIWeb** Arizona VII	**KEYS** Arizona VI	**MobiDE** Arizona I	**SWEET** Arizona V	**SWIM** Arizona III	**WebDB** Arizona IV
18:00- 20:00	**PODS Welcome Reception** Arizona I-V and Desert Garden								

MONDAY			
08:00- 08.30	**Continental Breakfast** Arizona, Vaquero, and South Foyers		
08:30- 09:45	**PODS Opening and Keynote Address** Arizona I-VIII		
09:45- 10:15	**Coffee Break** Arizona, South and West Foyers		
10:15- 11:45	**PODS 1** Streaming Arizona V-VIII	**DaMoN** Arizona I-III	**DBTest** Arizona IV
11:45- 13:15	**Lunch Provided for DaMoN and DBTest Participants** Vaquero E-F		
13:15- 14:30	**PODS 2** Awards Session Arizona V-VIII	**DaMoN** Arizona I-III	**DBTest** Arizona IV
14:30- 14:45	**Coffee Break** Arizona, South and West Foyers		
14:45- 15:45	**PODS 3** Tutorial Session 1 Arizona V-VIII	**DaMoN** Arizona I-III	**DBTest** Arizona IV
15:45- 16:15	**Coffee Break** Arizona, South and West Foyers		
16:15- 18:15	**PODS 4** Privacy and Semantic Web Arizona V-VIII	**DaMoN** Arizona I-III	**DBTest** Arizona IV
18:30- 19:30	**PODS Business Meeting** Arizona I-III		
19:00- 21:30	**SIGMOD Welcome Reception** **SIGMOD Undergraduate Research Poster Competition** **SIGMOD Informal Sponsor/Student Event** Terrace Court		

TUESDAY

08:00-08:30	**Continental Breakfast** Arizona, Vaquero, and South Foyers						
08:30-10:00	**SIGMOD Keynote Talk 1** Pat Hanrahan Arizona I-VIII						
10:00-10:30	**Coffee Break** Arizona, Vaquero, South, and West Foyers						
10:30-12:00	**PODS 5** Range Queries Arizona I-III	**Research 1** Distributed and Parallel Databases Arizona V	**Research 2** Indexing and Physical Database Design I Arroyo CDE	**Industry 1** Databases in the Cloud Arizona VI-VIII	**Demo A** Information Extraction, Search, Performance, and Clouds Vaquero A	**Tutorial 1** Mob Data Sourcing Arizona IV	
12:00-13:30	**Lunch by SAP** Vaquero D-G						
13:30-15:00	**PODS 6** Tutorial Session 2 Arizona II-III	**Research 3** Data Cleaning and Integration Arizona V	**Research 4** Query Processing and Optimization Arroyo CDE	**Industry 2** Social Media and Crowdsourcing Arizona VI-VIII	**Demo B** Social- or User-Centered Vaquero A	**Tutorial 2** Managing and Mining Large Graphs: Patterns and Algorithms Arizona IV	**Programming Contest** Arizona I
15:00-16:30	**PODS/SIGMOD Research Plenary Poster Session + Coffee** Vaquero B-C						
16:30-16:45	**Coffee Break** Arizona, Vaquero, South, and West Foyers						
16:45-18:15	**PODS 7** Views and Data Exchange Arizona I-III	**Information Session on Careers in Industry** Vaquero E-G		**Industry 3** Modern RDBMSs Arizona VI-VIII	**Demo C** Analytics Vaquero A	**Tutorial 3** Managing and Mining Large Graphs: Systems and Implementations Arizona IV	
18:15-18:30	**Coffee Break** West Foyer						
18:30-21:00	**SIGMOD New Researcher Symposium** Arizona I-IV						
21:00-	**Microsoft Reception** Arizona V-VIII						

WEDNESDAY						
08:00-08:30	**Continental Breakfast** Arizona, Vaquero, and South Foyers					
08:30-10:00	**SIGMOD Keynote Talk 2** Amin Vahdat Arizona I-VIII					
10:00-10:30	**Coffee Break** Arizona, Vaquero, South, and West Foyers					
10:30-12:00	**PODS 8** Indexing Arizona I-III	**Research 5** Social Networks and Graph Databases I Vaquero A	**Research 6** Data Visualization, Error Reporting Arizona V	**Research 7** Storage Systems, Query Processing and Optimization Arroyo CDE	**Industry 4** Big Data Arizona VI-VIII	**Tutorial 4** Computational Reproducibility: State-of-the-Art, Challenges, and Database Research Opportunities Arizona IV
12:00-14:00	**Lunch (provided by conference)** **SIGMOD Business Meeting** Vaquero D-G					
14:00-15:30	**PODS 9** Query Languages Arizona I-III	**SIGMOD Plenary Session: Poster Session for Workshop Papers** Vaquero B-C				
15:30-16:00	**Coffee Break** Arizona, South, and West Foyers					
16:00-17:30	**PODS 10** Streaming and Aggregation Arizona I-III	**SIGMOD Plenary Session: Perspectives on Big Data** Donald Kossmann, Kristen LeFevre, Sam Madden, Anand Rajaraman Arizona IV-VIII				
17:30-18:00	**PODS 10 (cont.)**					
18:00-18:30	**Departure to banquet site** Buses start departing at 18:00					
18:30-22:30	**Conference Banquet** Desert Foothills					

THURSDAY

08:00-08:30	**Continental Breakfast** Arizona, Vaquero, and South Foyers					
08:30-10:00	**SIGMOD Award Talks** Innovations, 10-Year, Dissertation, announcement of Best Demo Award winner Arizona I-VIII					
10:00-10:30	**Coffee Break** Arizona, Vaquero, South, and West Foyers					
10:30-12:00	**Research 8** Data Streams and Sensor Networks Arizona I-III	**Research 9** Mobile Databases Arizona V	**Research 10** Data Analytics Vaquero E	**Industry 5** Data Integration and Analytics Arizona VI-VIII	**Demo B** Social- or User-Centered Vaquero A	**Tutorial 5** Database Techniques for Linked Data Management Arizona IV
12:00-13:30	**Lunch (not provided by conference)** **Student/Researcher Vis-à-Vis Meeting (by invitation only, Vaquero F-G)**					
13:30-15:00	**Research 11** Crowdsourcing, Uncertainty in Databases Arizona I-III	**Research 12** Top-k Query Processing and Optimization Arizona V	**Research 13** Temporal and Graph Databases Vaquero E	**Industry 6** Query Processing and War Stories Arizona VI-VIII	**Demo C** Analytics Vaquero A	**Tutorial 6** Differential Privacy in Data Publication and Analysis Arizona IV
15:00-16:30	**SIGMOD Research Plenary Session + Coffee** Vaquero B-C					
16:30-16:45	**Coffee Break** Arizona, Vaquero, South, and West Foyers					
16:45-18:15	**Research 14** Information Retrieval and Text Mining Arizona I-III	**Research 15** Social Networks and Graph Databases II Arizona V	**Research 16** Indexing and Physical Database Design II Arizona VI-VIII		**Demo A** Information Extraction, Search, Performance, and Clouds Vaquero A	**Tutorial 7** Differential Privacy in Data Publication and Analysis Arizona IV

SESSION CONTENTS

SUNDAY, 18:00 – 20:00

PODS Welcome Reception
Location: Arizona Ballroom I–V and South Foyer / Desert Garden

MONDAY, 8:30 – 9:45

PODS Opening and Keynote Address
Location: Arizona Ballroom I–VIII
Hashtag: #pods12 #keynote
Session Chair: Maurizio Lenzerini (University of Rome La Sapienza)

What Next? A Half-Dozen Data Management Research Goals for Big Data and the Cloud
Surajit Chaudhuri, Microsoft Research

MONDAY, 10:15 – 11:45

PODS Session 1: Streaming
Location: Arizona Ballroom V–VIII
Hashtag: #pods12 #R1Px
Session Chair: Milan Vojnovic (Microsoft Research, Cambridge)

Graph Sketches: Sparsification, Spanners, and Subgraphs
Kook Jin Ahn, University of Pennsylvania; Sudipto Guha, University of Pennsylvania; Andrew McGregor, University of Massachusetts, Amherst

Approximating and Testing k-Histogram Distributions in Sub-linear Time
Piotr Indyk, MIT; Reut Levi, Tel Aviv University; Ronitt Rubinfeld, MIT

Mergeable Summaries
Pankaj Agarwal, Duke University; Graham Cormode, AT&T Labs-Research; Zengfeng Huang, The Hong Kong University of Science and Technology; Jeff Phillips, University of Utah; Zhewei Wei, The Hong Kong University of Science and Technology; Ke Yi, The Hong Kong University of Science and Technology

MONDAY, 13:15 – 14:30

PODS Session 2: Awards Session
Location: Arizona Ballroom V–VIII
Hashtag: #pods12 #R2Px
Session Chair: Richard Hull (IBM T. J. Watson Research Center)

The ACM PODS Alberto O. Mendelzon Test-of-Time Award: Containment and Equivalence for an XPath Fragment
Gerome Miklau, University of Massachusetts, Amherst; Dan Suciu, University of Washington

Best Paper Award: Worst-case Optimal Join Algorithms
Hung Q. Ngo, University at Buffalo, SUNY; Ely Porat, Bar-Ilan University; Christopher Ré, University of Wisconsin-Madison; Atri Rudra, University at Buffalo, SUNY

Regular Paper: Deterministic Regular Expressions in Linear Time
Benoît Groz, INRIA and University of Lille; Sebastian Maneth, NICTA and UNSW; Slawek Staworko, INRIA and University of Lille

MONDAY, 14:45 – 15:45

PODS Session 3: Tutorial Session 1
Location: Arizona Ballroom V–VIII
Hashtag: #pods12 #R3P1
Session Chair: Phokion Kolaitis (University of California, Santa Cruz and IBM Almaden Research Center)

Tutorial 1: Linguistic Foundations for Bidirectional Transformations
Benjamin C. Pierce, University of Pennsylvania

MONDAY, 16:15 – 18:15

PODS Session 4: Privacy and Semantic Web
Location: Arizona Ballroom V–VIII
Hashtag: #pods12 #R4Px
Session Chair: Pierre Senellart (Télécom ParisTech)

The Power of the Dinur-Nissim Algorithm: Breaking Privacy of Statistical and Graph Databases
Krzysztof Choromanski, Columbia University; Tal Malkin, Columbia University

A Rigorous and Customizable Framework for Privacy
Daniel Kifer, Penn State University; Ashwin Machanavajjhala, Yahoo! Research

Static Analysis and Optimization of Semantic Web Queries
Andrés Letelier, PUC Chile; Jorge Pérez, Universidad de Chile; Reinhard Pichler, Technische Universität Wien; Sebastian Skritek, Technische Universität Wien

The Complexity of Evaluating Path Expressions in SPARQL
Katja Losemann, Universität Bayreuth; Wim Martens, Universität Bayreuth

MONDAY, 18:30 – 19:30

PODS Business Meeting
Location: Arizona Ballroom I–III

MONDAY, 19:00 – 21:30

SIGMOD Welcome Reception
Location: Terrace Court

SIGMOD Informal Sponsor/Student event
Location: Terrace Court (in parallel with SIGMOD Reception)

SIGMOD Undergraduate Research Poster Competition
Location: Terrace Court (in parallel with SIGMOD Reception)
Hashtag: #sigmod12 #urpc

Declarative Web Application Development: Encapsulating Dynamic JavaScript Widgets
Robert Bolton, University of California, San Diego; David Ing, University of California, San Diego; Christopher Rebert, University of California, San Diego; Kristina Lam Thai, University of California, San Diego

Towards Scalable Summarization and Visualization of Large Text Corpora
Tyler Sliwkanich, University of Alberta; Douglas Schneider, University of Alberta; Aaron Yong, University of Alberta; Mitchell Home, University of Alberta; Denilson Barbosa, University of Alberta

Reducing Cache Misses in Hash Join Probing Phase By Pre-Sorting Strategy
Gi-Hwan Oh, SungKyunKwan University; Jae-Myung Kim, SungKyunKwan University; Woon-Hak Kang, SungKyunKwan University; Sang-Won Lee, SungKyunKwan University

DP-tree: Indexing Multi-Dimensional Data under Differential Privacy
Shangfu Peng, Shanghai Jiao Tong University; Yin Yang, Advanced Digital Sciences Center; Zhenjie Zhang, Advanced Digital Sciences Center; Marianne Winslett, Advanced Digital Sciences Center; Yong Yu, Shanghai Jiao Tong University

Temporal Provenance Discovery in Micro-Blog Message Streams
Zijun Xue, Peking University; Junjie Yao, Peking University; Bin Cui, Peking University

SigSpot: Mining Significant Anomalous Regions from Time-Evolving Networks
Misael Mongiovì, University of California, Santa Barbara; Petko Bogdanov, University of California, Santa Barbara; Razvan Ranca, University of California, Santa Barbara; Ambuj K. Singh, University of California, Santa Barbara; Evangelos E. Papalexakis, Carnegie Mellon University; Christos Faloutsos, Carnegie Mellon University

VRRC: Web Based Tool for Visualization and Recommendation on Co-Authorship Network
Eduardo M. Barbosa, UFMG; Mirella M. Moro, UFMG; Giseli Rabello Lopes, UFRGS; J. Palazzo M. de Oliveira, UFRGS

Fast Sampling Word Correlations of High Dimensional Text Data
Frank Rosner, Martin-Luther-University Halle-Wittenberg; Alexander Hinneburg, Martin-Luther-University Halle-Wittenberg; Martin Gleditzsch, Unister GmbH; Mathias Priebe, Unister GmbH; Andreas Both, Unister GmbH

TUESDAY, 8:30 – 10:00

SIGMOD Keynote Talk 1
Location: Arizona Ballroom I–VIII
Hashtag: #sigmod12 #keynote1
Session Chair: Luis Gravano (Columbia University)

Analytic Database Technologies for a New Kind of User - The Data Enthusiast
Pat Hanrahan, Stanford University and Tableau Software

TUESDAY, 10:30 – 12:00

PODS Session 5: Range Queries
Location: Arizona Ballroom I–III
Hashtag: #pods12 #R5Px
Session Chair: Srikanta Tirthapura (Iowa State)

Space-Efficient Range Reporting for Categorical Data
Yakov Nekrich, Universidad de Chile

Dynamic Top-K Range Reporting in External Memory
Cheng Sheng, Chinese University of Hong Kong; Yufei Tao, Korea Advanced Institute of Science and Technology

Indexability of 2D Range Search Revisited: Constant Redundancy and Weak Indivisibility
Yufei Tao, Chinese University of Hong Kong

SIGMOD Research 1: Distributed and Parallel Databases
Location: Arizona Ballroom V
Hashtag: #sigmod12 #R1Px
Session Chair: Anastasia Ailamaki (EPFL)

Calvin: Fast Distributed Transactions for Partitioned Database Systems
Alexander Thomson, Yale University; Thaddeus Diamond, Yale University; Shu-Chun Weng, Yale University; Kun Ren, Yale University; Philip Shao, Yale University; Daniel J. Abadi, Yale University

Advanced Partitioning Techniques for Massively Distributed Computation
Jingren Zhou, Microsoft; Nicolás Bruno, Microsoft; Wei Lin, Microsoft

SkewTune: Mitigating Skew in MapReduce Applications
YongChul Kwon, University of Washington; Magdalena Balazinska, University of Washington; Bill Howe, University of Washington; Jerome Rolia, HP Labs

SIGMOD Research 2: Indexing and Physical Database Design I
Location: Arroyo CDE
Hashtag: #sigmod12 #R2Px
Session Chair: Ashraf Aboulnaga (University of Waterloo)

Parallel Main-Memory Indexing for Moving-Object Query and Update Workloads
Darius Sidlauskas, Aalborg University; Simonas Saltenis, Aalborg University; Christian S. Jensen, Aarhus University

Divergent Physical Design Tuning for Replicated Databases
Mariano P. Consens, University of Toronto; Kleoni Ioannidou, University of California, Santa Cruz; Jeff LeFevre, University of California, Santa Cruz; Neoklis Polyzotis, University of California, Santa Cruz

Skew-Aware Automatic Database Partitioning in Shared-Nothing, Parallel OLTP Systems
Andrew Pavlo, Brown University; Carlo Curino, Yahoo! Research; Stanley Zdonik, Brown University

SIGMOD Industry 1: Databases in the Cloud
Location: Arizona Ballroom VI–VIII
Hashtag: #sigmod12 #I1Px
Session Chair: Jayant Madhavan (Google Research)

Amazon DynamoDB: A Seamlessly Scalable Non-Relational Datastore
Swami Sivasubramanian, Amazon

Efficient Transaction Processing in SAP HANA Database--The End of a Column Store Myth
Vishal Sikka, SAP; Franz Färber, SAP; Wolfgang Lehner, TUD/SAP; Sang Kyun Cha, SAP; Thomas Peh, SAP; Christof Bornhövd, SAP

Walnut: A Unified Cloud Object Store
Jianjun Chen, Yahoo!; Chris Douglas, Yahoo!; Michi Mutsuzaki, Yahoo!; Patrick Quaid, Yahoo!; Raghu Ramakrishnan, Yahoo!; Sriram Rao, Yahoo!; Russell Sears, Yahoo!

SIGMOD Demonstrations A: Information Extraction, Search, Performance, and Clouds
Location: Vaquero Ballroom A
Hashtag: #sigmod12 #demoAx

Automatic Web-Scale Information Extraction
Philip Bohannon, Yahoo! Research; Nilesh Dalvi, Yahoo! Research; Yuval Filmus, University of Toronto; Nori Jacoby, Yahoo!; Sathiya Keerthi, Yahoo! Research; Alok Kirpal, Yahoo! Research

Just-in-Time Information Extraction using Extraction Views
Amr El-Helw, EMC Corp.; Mina Farid, University of Waterloo; Ihab Ilyas, Qatar Computing Research Institute

ColumbuScout: Towards Building Local Search Engines over Large Databases
Cody Hansen, University of Utah; Feifei Li, University of Utah

Sofia Search: A Tool for Automating Related-Work Search
Behzad Golshan, Boston University; Theodoros Lappas, Boston University; Evimaria Terzi, Boston University

RACE: Real-Time Applications over Cloud-Edge
Badrish Chandramouli, Microsoft Research; Joris Claessens, Microsoft Research; Suman Nath, Microsoft Research; Ivo Santos, Microsoft Research; Wenchao Zhou, University of Pennsylvania

Partiqle: An Elastic SQL Engine over Key-Value Stores
Junichi Tatemura, NEC Laboratories America; Oliver Po, NEC Laboratories America; Wang-Pin Hsiung, NEC Laboratories America; Hakan Hacigümüs, NEC Laboratories America

JustMyFriends: Full SQL, Full Transactional Amenities, and Access Privacy
Arthur Meacham, New York University; Dennis Shasha, New York University

Dynamic Optimization of Generalized SQL Queries with Horizontal Aggregations
Carlos Ordonez, University of Houston; Javier García-García, UNAM; Zhibo Chen, University of Houston

ConsAD: A Real-Time Consistency Anomalies Detector
Kamal Zellag, McGill University; Bettina Kemme, McGill University

Interactive Performance Monitoring of a Composite OLTP and OLAP Workload
Anja Bog, Hasso Plattner Institute, University of Potsdam; Kai Sachs, SAP AG; Hasso Plattner, Hasso Plattner Institute, University of Potsdam

SIGMOD Tutorial 1: Mob Data Sourcing
Location: Arizona Ballroom IV
Hashtag: #sigmod12 #tutorial1
Presenters: Daniel Deutch, Ben Gurion University; Tova Milo, Tel Aviv University

TUESDAY, 13:30 – 15:00

PODS Session 6: Tutorial Session 2
Location: Arizona Ballroom II–III
Hashtag: #pods12 #R6Px
Session Chair: David Woodruff (IBM Almaden Research Center)

Tutorial 2: Approximate Computation and Implicit Regularization for Very Large-scale Data Analysis
Michael W. Mahoney, Stanford University

Regular Paper: Max-Sum Diversification, Monotone Submodular Functions and Dynamic Updates
Allan Borodin, University of Toronto; Hyun Chul Lee, LinkedIn Corporation; Yuli Ye, University of Toronto

SIGMOD Research 3: Data Cleaning and Integration
Location: Arizona Ballroom V
Hashtag: #sigmod12 #R3Px
Session Chair: Cong Yu (Google Research)

Sample-Driven Schema Mapping
Li Qian, University of Michigan, Ann Arbor; Michael J. Cafarella, University of Michigan, Ann Arbor; H. V. Jagadish, University of Michigan, Ann Arbor

Can we Beat the Prefix Filtering? An Adaptive Framework for Similarity Join and Search
Jiannan Wang, Tsinghua University; Guoliang Li, Tsinghua University; Jianhua Feng, Tsinghua University

InfoGather: Entity Augmentation and Attribute Discovery by Holistic Matching with Web Tables
Mohamed Yakout, Purdue University; Kris Ganjam, Microsoft Research; Kaushik Chakrabarti, Microsoft Research; Surajit Chaudhuri, Microsoft Research

SIGMOD Research 4: Query Processing and Optimization
Location: Arroyo CDE
Hashtag: #sigmod12 #R4Px
Session Chair: Anish Das Sarma (Google Research)

Interactive Regret Minimization
Danupon Nanongkai, University of Vienna; Ashwin Lall, Denison University; Atish Das Sarma, Google Research; Kazuhisa Makino, University of Tokyo

MCJoin: A Memory-Constrained Join for Column-Store Main-Memory Databases
Steven Keith Begley, La Trobe University; Zhen He, La Trobe University; Yi-Ping Phoebe Chen, La Trobe University

Holistic Optimization by Prefetching Query Results
Karthik Ramachandra, Indian Institute of Technology Bombay; S. Sudarshan, Indian Institute of Technology Bombay

SIGMOD Industry 2: Social Media and Crowdsourcing
Location: Arizona Ballroom VI–VIII
Hashtag: #sigmod12 #I2Px
Session Chair: Zachary Ives (University of Pennsylvania)

The Value of Social Media Data in Enterprise Applications
Shivakumar Vaithyanathan, IBM Almaden Research Center

Anatomy of a Gift Recommendation Engine Powered by Social Media
Yannis Pavlidis, @WalmartLabs; Madhusudan Mathihalli, @WalmartLabs; Indrani Chakravarty, @WalmartLabs; Arvind Batra, @WalmartLabs; Ron Benson, @WalmartLabs; Ravi Raj, @WalmartLabs; Robert Yau, @WalmartLabs; Mike McKiernan, @WalmartLabs; Venky Harinarayan, @WalmartLabs; Anand Rajaraman, @WalmartLabs

Designing a Scalable Crowdsourcing Platform
Chris Van Pelt, CrowdFlower; Alex Sorokin, CrowdFlower

SIGMOD Demonstrations B: Social- or User-Centered
Location: Vaquero Ballroom A
Hashtag: #sigmod12 #demoBx

Sindbad: A Location-Based Social Networking System
Mohamed Sarwat, University of Minnesota; Jie Bao, University of Minnesota; Ahmed Eldawy, University Of Minnesota; Justin Levandoski, Microsoft Research; Amr Magdy, University of Minnesota; Mohamed Mokbel, University of Minnesota

MAQSA: A System for Social Analytics on News
Sihem Amer-Yahia, Qatar Computing Research Institute; Samreen Anjum, Qatar Computing Research Institute; Amira Ghenai, Qatar Computing Research Institute; Aysha Siddique, Qatar Computing Research Institute; Sofiane Abbar, Qatar Computing Research Institute; Sam Madden, MIT; Adam Marcus, MIT; Mohammed El-Haddad; Al Jazeera Network

Surfacing Time-Critical Insights from Social Media

Bogdan Alexe, IBM Almaden Research Center; Mauricio Hernandez, IBM Almaden Research Center; Kirsten Hildrum, IBM T. J. Watson Research Center; Rajasekar Krishnamurthy, IBM Almaden Research Center; Georgia Koutrika, IBM Almaden Research Center; Meenakshi Nagarajan, IBM Almaden Research Center; Haggai Roitman, IBM Research, Haifa; Michal Shmueli-Scheuer, IBM Research, Haifa; Ioana Stanoi, IBM Almaden Research Center; Chitra Venkatramani, IBM T. J. Watson Research Center; Rohit Wagle, IBM T. J. Watson Research Center

Taagle: Efficient, Personalized Search in Collaborative Tagging Networks

Silviu Maniu, Télécom ParisTech, CNRS LTCI; Bogdan Cautis, Télécom ParisTech, CNRS LTCI

PrefDB: Bringing Preferences Closer to the DBMS

Anastasios Arvanitis, National Technical University of Athens; Georgia Koutrika, IBM Almaden Research Center

Auto-Completion Learning for XML

Serge Abiteboul, Collège de France, INRIA Saclay, ENS Cachan; Yael Amsterdamer, Tel Aviv University; Tova Milo, Tel Aviv University; Pierre Senellart, Télécom ParisTech, CNRS LTCI

Logos: A System for Translating Queries into Narratives

Andreas Kokkalis, University of Athens; Panagiotis Vagenas, University of Athens; Alexandros Zervakis, University of Athens; Alkis Simitsis, HP Labs; Georgia Koutrika, IBM Almaden Research Center; Yannis Ioannidis, University of Athens

PAnG-Finding Patterns in Annotation Graphs

Philip Anderson, University of Maryland; Andreas Thor, University of Maryland; Joseph Benik, University of Maryland; Louiqa Raschid, University of Maryland; María Esther Vidal, Universidad Simón Bolívar

VizDeck: Self-Organizing Dashboards for Visual Analytics

Alicia Key, University of Washington; Bill Howe, University of Washington; Daniel Perry, University of Washington; Cecilia Aragon, University of Washington

Kaizen: A Semi-Automatic Index Advisor

Ivo Jimenez, University of California, Santa Cruz; Huascar Sanchez, University of California, Santa Cruz; Quoc Trung Tran, University of California, Santa Cruz; Neoklis Polyzotis, University of California, Santa Cruz

SIGMOD Tutorial 2: Managing and Mining Large Graphs: Patterns and Algorithms
Location: Arizona Ballroom IV
Hashtag: #sigmod12 #tutorial2
Presenters: Christos Faloutsos, Carnegie Mellon University; U Kang, Carnegie Mellon University

SIGMOD Programming Contest: Finalist Presentations and Panel
Location: Arizona Ballroom I
Hashtag: #sigmod12 #pc
Session Chair: Thomas Kissinger (TU Dresden)

TUESDAY, 15:00 – 16:30

PODS/SIGMOD Research Plenary Poster Session
Location: Vaquero Ballroom B–C
Hashtag: #sigmod12 #pods12 #rp1

Papers from all PODS Sessions and from SIGMOD Sessions Research 1 to 7

TUESDAY, 16:45 – 18:15

PODS Session 7: Views and Data Exchange
Location: Arizona Ballroom I–III
Hashtag: #pods12 #R7Px
Session Chair: Jan Paredaens (University of Antwerp)

Query-Based Data Pricing
Paraschos Koutris, University of Washington; Prasang Upadhyaya, University of Washington; Magdalena Balazinska, University of Washington; Bill Howe, University of Washington; Dan Suciu, University of Washington

Local Transformations and Conjunctive-Query Equivalence
Ronald Fagin, IBM Almaden Research Center; Phokion G. Kolaitis, University of California, Santa Cruz and IBM Almaden Research Center

A Dichotomy in the Complexity of Deletion Propagation with Functional Dependencies
Benny Kimelfeld, IBM Almaden Research Center

SIGMOD Information Session on Careers in Industry
Location: Vaquero Ballroom E–G
Hashtag: #sigmod12 #career

SIGMOD Industry 3: Modern RDBMSs
Location: Arizona Ballroom VI–VIII
Hashtag: #sigmod12 #I3Px
Session Chair: Rachel Pottinger (University of British Columbia)

Query Optimization in Microsoft SQL Server PDW
Srinath Shankar, Microsoft; Rimma Nehme, Microsoft; Josep Aguilar-Saborit, Microsoft; Andrew Chung, Microsoft; Mostafa Elhemali, Microsoft; Alan Halverson, Microsoft; Eric Robinson, Microsoft; Mahadevan Sankara Subramanian, Microsoft; David DeWitt, Microsoft; César Galindo-Legaria, Microsoft

F1—The Fault-Tolerant Distributed RDBMS Supporting Google's Ad Business
Jeff Shute, Google; Mircea Oancea, Google; Stephan Ellner, Google; Ben Handy, Google; Eric Rollins, Google; Bart Samwel, Google; Radek Vingralek, Google; Chad Whipkey, Google; Xin Chen, Google; Beat Jegerlehner, Google; Kyle Littlefield, Google; Phoenix Tong, Google

Oracle In-Database Hadoop: When MapReduce Meets RDBMS
Xueyuan Su, Yale University; Garret Swart, Oracle

SIGMOD Demonstrations C: Analytics
Location: Vaquero Ballroom A
Hashtag: #sigmod12 #demoCx

Shark: Fast Data Analysis Using Coarse-Grained Distributed Memory
Cliff Engle, University of California, Berkeley; Antonio Lupher, University of California, Berkeley; Reynold Xin, University of California, Berkeley; Matei Zaharia, University of California, Berkeley; Michael Franklin, University of California, Berkeley; Scott Shenker, University of California, Berkeley; Ion Stoica, University of California, Berkeley

Exploiting MapReduce-Based Similarity Joins
Yasin N. Silva, Arizona State University; Jason M. Reed, Arizona State University

GLADE: Big Data Analytics Made Easy
Yu Cheng, University of California, Merced; Chengjie Qin, University of California, Merced; Florin Rusu, University of California, Merced

ReStore: Reusing Results of MapReduce Jobs in Pig
Iman Elghandour, University of Waterloo; Ashraf Aboulnaga, University of Waterloo

Clydesdale: Structured Data Processing on Hadoop

Andrey Balmin, IBM Almaden Research Center; Tim Kaldewey, IBM Almaden Research Center; Sandeep Tata, IBM Almaden Research Center

Tiresias: A Demonstration of How-To Queries

Alexandra Meliou, University of Washington; Yisong Song, University of Washington; Dan Suciu, University of Washington

AstroShelf: Understanding the Universe Through Scalable Navigation of a Galaxy of Annotations

Panayiotis Neophytou, University of Pittsburgh; Roxana Gheorghiu, University of Pittsburgh; Rebecca Hachey, University of Pittsburgh; Timothy Luciani, University of Pittsburgh; Di Bao, University of Pittsburgh; Alexandros Labrinidis, University of Pittsburgh; Elisabeta G. Marai, University of Pittsburgh; Panos K. Chrysanthis, University of Pittsburgh

OPAvion: Mining and Visualization in Large Graphs

Leman Akoglu, Carnegie Mellon University; Duen Horng Chau, Carnegie Mellon University; U Kang, Carnegie Mellon University; Danai Koutra, Carnegie Mellon University; Christos Faloutsos, Carnegie Mellon University

CloudAlloc: A Monitoring and Reservation System for Compute Clusters

Enrico Iori, University of Trento; Alkis Simitsis, HP Labs; Themis Palpanas, University of Trento; Kevin Wilkinson, HP Labs; Stavros Harizopoulos, Nou Data

TIRAMOLA: Elastic NoSQL Provisioning through a Cloud Management Platform

Ioannis Konstantinou, National Technical University of Athens; Evangelos Angelou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Christina Boumpouka, National Technical University of Athens; Nectarios Koziris, National Technical University of Athens; Spyros Sioutas, Ionian University

SIGMOD Tutorial 3: Managing and Mining Large Graphs: Systems and Implementations

Location: Arizona Ballroom IV

Hashtag: #sigmod12 #tutorial3

Presenters: Bin Shao, Microsoft Research Asia; Haixun Wang, Microsoft Research Asia; Yanhua Xiao, Fudan University and Microsoft Research Asia

TUESDAY, 18:30 – 21:00

SIGMOD New Researcher Symposium: "How to be a good advisor/advisee?"

Location: Arizona Ballroom I–IV

Hashtag: #sigmod12 #nrs

Session Chairs: Xin Luna Dong (AT&T Labs-Research) and Christopher Ré (University of Wisconsin-Madison)

Session 1

Magdalena Balazinska, University of Washington; Yi Chen, Arizona State University; Anish Das Sarma, Google Research; Lukasz Golab, University of Waterloo

Session 2

Sihem Amer-Yahia, QCRI; Christian S. Jensen, Aarhus University; Henry Korth, Lehigh University; M. Tamer Özsu, University of Waterloo

TUESDAY, 21:00 –

Microsoft Reception
Location: Arizona Ballroom V–VIII

WEDNESDAY, 8:30 – 10:00

SIGMOD Keynote Talk 2
Location: Arizona Ballroom I–VIII
Hashtag: #sigmod12 #keynote2
Session Chair: Surajit Chaudhuri (Microsoft Research)

Symbiosis in Scale Out Networking and Data Management
Amin Vahdat, University of California San Diego and Google

WEDNESDAY, 10:30 – 12:00

PODS Session 8: Indexing
Location: Arizona Ballroom I–III
Hashtag: #pods12 #R8Px
Session Chair: Yakov Nekrich (University of Bonn)

The Wavelet Trie: Maintaining an Indexed Sequence of Strings in Compressed Space
Roberto Grossi, Università di Pisa; Giuseppe Ottaviano, Università di Pisa

On the Optimality of Clustering Properties of Space Filling Curves
Pan Xu, Iowa State University; Srikanta Tirthapura, Iowa State University

Nearest-Neighbor Searching Under Uncertainty
Pankaj K. Agarwal, Duke University; Alon Efrat, The University of Arizona; Swaminathan Sankararaman, Duke University; Wuzhou Zhang, Duke University

SIGMOD Research 5: Social Networks and Graph Databases I
Location: Vaquero Ballroom A
Hashtag: #sigmod12 #R5Px
Session Chair: Minos Garofalakis (Technical University of Crete)

Managing Large Dynamic Graphs Efficiently
Jayanta Mondal, University of Maryland; Amol Deshpande, University of Maryland

Query Preserving Graph Compression
Wenfei Fan, University of Edinburgh; Jianzhong Li, Harbin Institute of Technology; Xin Wang, University of Edinburgh; Yinghui Wu, University of Edinburgh and University of California, Santa Barbara

SCARAB: Scaling Reachability Computation on Large Graphs
Ruoming Jin, Kent State University; Ning Ruan, Kent State University; Saikat Dey, Kent State University; Jeffrey Xu Yu, The Chinese University of Hong Kong

SIGMOD Research 6: Data Visualization, Error Reporting
Location: Arizona Ballroom V
Hashtag: #sigmod12 #R6Px
Session Chair: Yannis Papakonstantinou (University of California, San Diego)

Skimmer: Rapid Scrolling of Relational Query Results
Manish Singh, University of Michigan, Ann Arbor; Arnab Nandi, Ohio State University; H. V. Jagadish, University of Michigan, Ann Arbor

Efficient Spatial Sampling of Large Geographical Tables
Anish Das Sarma, Google; Hongrae Lee, Google; Hector Gonzalez, Google; Jayant Madhavan, Google; Alon Halevy, Google

Declarative Error Management for Robust Data-Intensive Applications
Carl-Christian Kanne, Platfora Inc.; Vuk Ercegovac, IBM Almaden Research Center

SIGMOD Research 7: Storage Systems, Query Processing and Optimization
Location: Arroyo CDE
Hashtag: #sigmod12 #R7Px
Session Chair: Ioana Manolescu (INRIA)

bLSM: A General Purpose Log Structured Merge Tree
Russell Sears, Yahoo! Research; Raghu Ramakrishnan, Yahoo! Research

Skeleton Automata for FPGAs: Reconfiguring without Reconstructing
Jens Teubner, ETH Zürich; Louis Woods, ETH Zürich; Chongling Nie, ETH Zürich

NoDB: Efficient Query Execution on Raw Data Files
Ioannis Alagiannis, EPFL; Renata Borovica, EPFL; Miguel Branco, EPFL; Stratos Idreos, CWI; Anastasia Ailamaki, EPFL

SIGMOD Industry 4: Big Data
Location: Arizona Ballroom VI–VIII
Hashtag: #sigmod12 #I4Px
Session Chair: Cong Yu (Google Research)

TAO: How Facebook Serves the Social Graph
Venkateshwaran Venkataramani, Facebook; Zach Amsden, Facebook; Nathan Bronson, Facebook; George Cabrera III, Facebook; Prasad Chakka, Facebook; Peter Dimov, Facebook; Hui Ding, Facebook; Jack Ferris, Facebook; Anthony Giardullo, Facebook; Jeremy Hoon, Facebook; Sachin Kulkarni, Facebook; Nathan Lawrence, Facebook; Mark Marchukov, Facebook; Dmitri Petrov, Facebook; Lovro Puzar, Facebook

Large-Scale Machine Learning at Twitter
Jimmy Lin, Twitter; Alex Kolcz, Twitter

Recurring Job Optimization in Scope
Nicolás Bruno, Microsoft; Sameer Agarwal, Microsoft; Srikanth Kandula, Microsoft; Bing Shi, Microsoft; Ming-Chuan Wu, Microsoft; Jingren Zhou, Microsoft

SIGMOD Tutorial 4: Computational Reproducibility: State-of-the-Art, Challenges, and Database Research Opportunities
Location: Arizona Ballroom IV
Hashtag: #sigmod12 #tutorial4
Presenters: Juliana Freire, NYU-Poly; Philippe Bonnet, IT University of Copenhagen; Dennis Shasha, NYU

WEDNESDAY, 12:00 – 14:00

SIGMOD Business Meeting
Location: Vaquero Ballroom D–G
Session Chair: Yannis Ioannidis (University of Athens)

WEDNESDAY, 14:00 – 15:30

PODS Session 9: Query Languages
Location: Arizona Ballroom I–III
Hashtag: #pods12 #R9Px
Session Chair: Daniel Kifer (Penn State University)

Classification of Annotation Semirings over Query Containment
Egor V. Kostylev, University of Edinburgh; Juan L. Reutter, University of Edinburgh; András Z. Salamon, University of Edinburgh

Efficient Approximations of Conjunctive Queries
Pablo Barceló, Universidad de Chile; Leonid Libkin, University of Edinburgh; Miguel Romero, Universidad de Chile

On the Complexity of Package Recommendation Problems
Ting Deng, Beihang University; Wenfei Fan, University of Edinburgh; Floris Geerts, University of Antwerp

SIGMOD Plenary Session: Poster Session for Workshop Papers
Location: Vaquero Ballroom B–C
Hashtag: #sigmod12 #rp2

WEDNESDAY, 16:00 – 17:30

PODS Session 10: Streaming and Aggregation
Location: Arizona Ballroom I–III
Hashtag: #pods12 #R10Px
Session Chair: Yufei Tao (Chinese University of Hong Kong)

Space-Efficient Estimation of Statistics over Sub-Sampled Streams
Andrew McGregor, University of Massachusetts, Amherst; A. Pavan, Iowa State University; Srikanta Tirthapura, Iowa State University; David Woodruff, IBM Almaden Research Center

Rectangle-Efficient Aggregation in Spatial Data Streams
Srikanta Tirthapura, Iowa State University; David Woodruff, IBM Almaden Research Center

Randomized Algorithms for Tracking Distributed Count, Frequencies, and Ranks
Zengfeng Huang, Hong Kong University of Science and Technology; Ke Yi, Hong Kong University of Science and Technology; Qin Zhang, Aarhus University

Continuous Distributed Counting for Non-monotonic Streams
Zhenming Liu, Harvard University; Božidar Radunović, Microsoft Research; Milan Vojnović, Microsoft Research

SIGMOD Plenary Session: Perspectives on Big Data
Location: Arizona Ballroom IV–VIII
Hashtag: #sigmod12 #bigdata
Session Chair: Surajit Chaudhuri (Microsoft Research)

Presenters: Donald Kossmann (ETHZ)
Kristen LeFevre (Google Research and University of Michigan)
Sam Madden (MIT)
Anand Rajaraman (@WalmartLabs)

WEDNESDAY, 17:30 – 18:00

PODS Session 10 (continued)
Location: Arizona Ballroom I–III

WEDNESDAY, 18:00 – 18:30

Departure to banquet site
Buses start departing at 18:00.

WEDNESDAY, 18:30 – 22:30

Conference Banquet
Buses start departing at 18:00.
Location: Desert Foothills

THURSDAY, 8:30 – 10:00

SIGMOD Award Talks (Innovations, 10-Year, Dissertation, announcement of Best Demo Award winner)
Location: Arizona Ballroom I–VIII
Hashtag: #sigmod12 #award
Session Chair: Rakesh Agrawal (Microsoft Research)

THURSDAY, 10:30 – 12:00

SIGMOD Research 8: Data Streams and Sensor Networks
Location: Arizona Ballroom I–III
Hashtag: #sigmod12 #R8Px
Session Chair: Henry Korth (Lehigh University)

Best Paper Award: High-Performance Complex Event Processing over XML Streams
Barzan Mozafari, University of California, Los Angeles; Kai Zeng, University of California, Los Angeles; Carlo Zaniolo, University of California, Los Angeles

Prediction-Based Geometric Monitoring over Distributed Data Streams
Nikos Giatrakos, University of Piraeus; Antonios Deligiannakis, Technical University of Crete; Minos Garofalakis, Technical University of Crete; Izchak Sharfman, Technion; Assaf Schuster, Technion

Online Windowed Subsequence Matching over Probabilistic Sequences
Zheng Li, University of Massachusetts, Lowell; Tingjian Ge, University of Massachusetts, Lowell

SIGMOD Research 9: Mobile Databases
Location: Arizona Ballroom V
Hashtag: #sigmod12 #R9Px
Session Chair: Ryan Johnson (University of Toronto)

MaskIt: Privately Releasing User Context Streams for Personalized Mobile Applications
Michaela Goetz, Twitter; Suman Nath, Microsoft Research; Johannes Gehrke, Cornell University

Authenticating Location-Based Services without Compromising Location Privacy
Haibo Hu, Hong Kong Baptist University; Jianliang Xu, Hong Kong Baptist University; Qian Chen, Hong Kong Baptist University; Ziwei Yang, Hong Kong Baptist University

Effective Caching of Shortest Paths for Location-Based Services
Jeppe Rishede Thomsen, Hong Kong Polytechnic University; Man Lung Yiu, Hong Kong Polytechnic University; Christian S. Jensen, Aarhus University

SIGMOD Research 10: Data Analytics
Location: Vaquero Ballroom E
Hashtag: #sigmod12 #R10Px
Session Chair: Stratos Idreos (CWI)

Towards a Unified Architecture for in-RDBMS Analytics
Xixuan Feng, University of Wisconsin-Madison; Arun Kumar, University of Wisconsin-Madison; Benjamin Recht, University of Wisconsin-Madison; Christopher Ré, University of Wisconsin-Madison

Tiresias: The Database Oracle for How-To Queries
Alexandra Meliou, University of Washington; Dan Suciu, University of Washington

GUPT: Privacy Preserving Data Analysis Made Easy
Prashanth Mohan, University of California, Berkeley; Abhradeep Thakurta, Pennsylvania State University; Elaine Shi, University of California, Berkeley; Dawn Song, University of California, Berkeley; David Culler, University of California, Berkeley

SIGMOD Industry 5: Data Integration and Analytics
Location: Arizona Ballroom VI–VIII
Hashtag: #sigmod12 #I5Px
Session Chair: Xin Luna Dong (AT&T Labs-Research)

Dynamic Workload-Driven Data Integration in Tableau
Kristi Morton, University of Washington; Ross Bunker, Tableau Software; Jock Mackinlay, Tableau Software; Robert Morton, Tableau Software; Chris Stolte, Tableau Software

Finding Related Tables
Anish Das Sarma, Google; Lujun Fang, Google; Nitin Gupta, Google; Alon Halevy, Google; Hongrae Lee, Google; Fei Wu, Google; Reynold Xin, Google; Cong Yu, Google

Optimizing Analytic Data Flows for Multiple Execution Engines
Alkis Simitsis, HP Labs; Kevin Wilkinson, HP Labs; Malu Castellanos, HP Labs; Umeshwar Dayal, HP Labs

SIGMOD Demonstrations B: Social- or User-Centered
Location: Vaquero Ballroom A
Hashtag: #sigmod12 #demoBx

Sindbad: A Location-Based Social Networking System

Mohamed Sarwat, University of Minnesota; Jie Bao, University of Minnesota; Ahmed Eldawy, University Of Minnesota; Justin Levandoski, Microsoft Research; Amr Magdy, University of Minnesota; Mohamed Mokbel, University of Minnesota

MAQSA: A System for Social Analytics on News

Sihem Amer-Yahia, Qatar Computing Research Institute; Samreen Anjum, Qatar Computing Research Institute; Amira Ghenai, Qatar Computing Research Institute; Aysha Siddique, Qatar Computing Research Institute; Sofiane Abbar, Qatar Computing Research Institute; Sam Madden, MIT; Adam Marcus, MIT; Mohammed El-Haddad; Al Jazeera Network

Surfacing Time-Critical Insights from Social Media

Bogdan Alexe, IBM Almaden Research Center; Mauricio Hernandez, IBM Almaden Research Center; Kirsten Hildrum, IBM T. J. Watson Research Center; Rajasekar Krishnamurthy, IBM Almaden Research Center; Georgia Koutrika, IBM Almaden Research Center; Meenakshi Nagarajan, IBM Almaden Research Center; Haggai Roitman, IBM Research, Haifa; Michal Shmueli-Scheuer, IBM Research, Haifa; Ioana Stanoi, IBM Almaden Research Center; Chitra Venkatramani, IBM T. J. Watson Research Center; Rohit Wagle, IBM T. J. Watson Research Center

Taagle: Efficient, Personalized Search in Collaborative Tagging Networks

Silviu Maniu, Télécom ParisTech, CNRS LTCI; Bogdan Cautis, Télécom ParisTech, CNRS LTCI

PrefDB: Bringing Preferences Closer to the DBMS

Anastasios Arvanitis, National Technical University of Athens; Georgia Koutrika, IBM Almaden Research Center

Auto-Completion Learning for XML

Serge Abiteboul, Collège de France, INRIA Saclay, ENS Cachan; Yael Amsterdamer, Tel Aviv University; Tova Milo, Tel Aviv University; Pierre Senellart, Télécom ParisTech, CNRS LTCI

Logos: A System for Translating Queries into Narratives

Andreas Kokkalis, University of Athens; Panagiotis Vagenas, University of Athens; Alexandros Zervakis, University of Athens; Alkis Simitsis, HP Labs; Georgia Koutrika, IBM Almaden Research Center; Yannis Ioannidis, University of Athens

PAnG-Finding Patterns in Annotation Graphs

Philip Anderson, University of Maryland; Andreas Thor, University of Maryland; Joseph Benik, University of Maryland; Louiqa Raschid, University of Maryland; María Esther Vidal, Universidad Simón Bolívar

VizDeck: Self-Organizing Dashboards for Visual Analytics

Alicia Key, University of Washington; Bill Howe, University of Washington; Daniel Perry, University of Washington; Cecilia Aragon, University of Washington

Kaizen: A Semi-Automatic Index Advisor

Ivo Jimenez, University of California, Santa Cruz; Huascar Sanchez, University of California, Santa Cruz; Quoc Trung Tran, University of California, Santa Cruz; Neoklis Polyzotis, University of California, Santa Cruz

SIGMOD Tutorial 5: Database Techniques for Linked Data Management
Location: Arizona Ballroom IV
Hashtag: #sigmod12 #tutorial5
Presenters: Andreas Harth, Karlsruhe Institute of Technology (KIT); Katja Hose, Max-Planck Institute for Informatics; Ralf Schenkel, Saarland University

THURSDAY, 12:00–13:30

Student/Researcher Vis-à-Vis Meeting (by invitation only)
Location: Vaquero Ballroom F–G

THURSDAY, 13:30 – 15:00

SIGMOD Research 11: Crowdsourcing, Uncertainty in Databases
Location: Arizona Ballroom I–III
Hashtag: #sigmod12 #R11Px
Session Chair: Martin Theobald (Max-Planck Institute for Informatics)

CrowdScreen: Algorithms for Filtering Data with Humans
Aditya G. Parameswaran, Stanford University; Héctor García-Molina, Stanford University; Hyunjung Park, Stanford University; Neoklis Polyzotis, University of California, Santa Cruz; Aditya Ramesh, Stanford University; Jennifer Widom, Stanford University

Local Structure and Determinism in Probabilistic Databases
Theodoros Rekatsinas, University of Maryland; Amol Deshpande, University of Maryland; Lise Getoor, University of Maryland

So Who Won? Dynamic Max Discovery with the Crowd
Stephen Guo, Stanford University; Aditya G. Parameswaran, Stanford University; Héctor García-Molina, Stanford University

SIGMOD Research 12: Top-k Query Processing and Optimization
Location: Arizona Ballroom V
Hashtag: #sigmod12 #R12Px
Session Chair: Vasilis Vassalos (Athens University of Economics and Business)

Processing a Large Number of Continuous Preference Top-k Queries
Albert Yu, Duke University; Pankaj K. Agarwal, Duke University; Jun Yang, Duke University

Optimal Top-k Generation of Attribute Combinations Based on Ranked Lists
Jiaheng Lu, Renmin University of China; Pierre Senellart, Télécom ParisTech; Chunbin Lin, Renmin University of China; Xiaoyong Du, Renmin University of China; Shan Wang, Renmin University of China; Xinxing Chen, Renmin University of China

Top-k Bounded Diversification
Piero Fraternali, Politecnico di Milano; Davide Martinenghi, Politecnico di Milano; Marco Tagliasacchi, Politecnico di Milano

SIGMOD Research 13: Temporal and Graph Databases
Location: Vaquero Ballroom E
Hashtag: #sigmod12 #R13Px
Session Chair: Philip Bohannon (Yahoo! Research)

Temporal Alignment
Anton Dignös, University of Zürich; Michael H. Böhlen, University of Zürich; Johann Gamper, Free University of Bozen-Bolzano

A Highway-Centric Labeling Approach for Answering Distance Queries on Large Sparse Graphs
Ruoming Jin, Kent State University; Ning Ruan, Kent State University; Yang Xiang, Ohio State University; Victor Lee, Kent State University

Efficient Processing of Distance Queries in Large Graphs: A Vertex Cover Approach
James Cheng, Nanyang Technological University; Yiping Ke, Institute of High Performance Computing, Singapore; Shumo Chu, Nanyang Technological University; Carter Cheng, Nanyang Technological University

SIGMOD Industry 6: Query Processing and War Stories
Location: Arizona Ballroom VI–VIII
Hashtag: #sigmod12 #I6Px
Session Chair: AnHai Doan (University of Wisconsin-Madison and @WalmartLabs)

CloudRAMSort: Fast and Efficient Large-Scale Distributed RAM Sort on Shared-Nothing Cluster

Changkyu Kim, Intel Labs; Jongsoo Park, Intel Labs; Nadathur Satish, Intel Labs; Hongrae Lee, Google Research; Pradeep Dubey, Intel Labs; Jatin Chhugani, Intel Labs

Adaptive Optimizations of Recursive Queries in Teradata

Ahmad Ghazal, Teradata; Dawit Seid, Teradata; Alain Crolotte, Teradata; Mohammed Al-Kateb, Teradata

From X100 to Vectorwise: Opportunities, Challenges and Things Most Researchers Do Not Think About

Marcin Zukowski, Actian; Peter Boncz, CWI

SIGMOD Demonstrations C: Analytics
Location: Vaquero Ballroom A
Hashtag: #sigmod12 #demoCx

Shark: Fast Data Analysis Using Coarse-Grained Distributed Memory

Cliff Engle, University of California, Berkeley; Antonio Lupher, University of California, Berkeley; Reynold Xin, University of California, Berkeley; Matei Zaharia, University of California, Berkeley; Michael Franklin, University of California, Berkeley; Scott Shenker, University of California, Berkeley; Ion Stoica, University of California, Berkeley

Exploiting MapReduce-Based Similarity Joins

Yasin N. Silva, Arizona State University; Jason M. Reed, Arizona State University

GLADE: Big Data Analytics Made Easy

Yu Cheng, University of California, Merced; Chengjie Qin, University of California, Merced; Florin Rusu, University of California, Merced

ReStore: Reusing Results of MapReduce Jobs in Pig

Iman Elghandour, University of Waterloo; Ashraf Aboulnaga, University of Waterloo

Clydesdale: Structured Data Processing on Hadoop

Andrey Balmin, IBM Almaden Research Center; Tim Kaldewey, IBM Almaden Research Center; Sandeep Tata, IBM Almaden Research Center

Tiresias: A Demonstration of How-To Queries

Alexandra Meliou, University of Washington; Yisong Song, University of Washington; Dan Suciu, University of Washington

AstroShelf: Understanding the Universe Through Scalable Navigation of a Galaxy of Annotations

Panayiotis Neophytou, University of Pittsburgh; Roxana Gheorghiu, University of Pittsburgh; Rebecca Hachey, University of Pittsburgh; Timothy Luciani, University of Pittsburgh; Di Bao, University of Pittsburgh; Alexandros Labrinidis, University of Pittsburgh; Elisabeta G. Marai, University of Pittsburgh; Panos K. Chrysanthis, University of Pittsburgh

OPAvion: Mining and Visualization in Large Graphs

Leman Akoglu, Carnegie Mellon University; Duen Horng Chau, Carnegie Mellon University; U Kang, Carnegie Mellon University; Danai Koutra, Carnegie Mellon University; Christos Faloutsos, Carnegie Mellon University

CloudAlloc: A Monitoring and Reservation System for Compute Clusters

Enrico Iori, University of Trento; Alkis Simitsis, HP Labs; Themis Palpanas, University of Trento; Kevin Wilkinson, HP Labs; Stavros Harizopoulos, Nou Data

TIRAMOLA: Elastic NoSQL Provisioning through a Cloud Management Platform

Ioannis Konstantinou, National Technical University of Athens; Evangelos Angelou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Christina Boumpouka, National Technical University of Athens; Nectarios Koziris, National Technical University of Athens; Spyros Sioutas, Ionian University

SIGMOD Tutorial 6: Differential Privacy in Data Publication and Analysis (Part 1)

Location: Arizona Ballroom IV
Hashtag: #sigmod12 #tutorial6
Presenters: Yin Yang, Advanced Digital Sciences Center, Singapore; Zhenjie Zhang, Advanced Digital Sciences Center, Singapore; Gerome Miklau, University of Massachusetts, Amherst; Marianne Winslett, University of Illinois at Urbana-Champaign; Xiaokui Xiao, Nanyang Technological University

THURSDAY, 15:00 – 16:30

SIGMOD Research Plenary Poster Session

Location: Vaquero Ballroom B–C
Hashtag: #sigmod12 #rp3

Papers from SIGMOD Sessions Research 8 to 16

THURSDAY, 16:45 – 18:15

SIGMOD Research 14: Information Retrieval and Text Mining

Location: Arizona Ballroom I–III
Hashtag: #sigmod12 #R14Px
Session Chair: Jun Yang (Duke University)

Aggregate Suppression for Enterprise Search Engines
Mingyang Zhang, George Washington University; Nan Zhang, George Washington University; Gautam Das, University of Texas at Arlington

Probase: A Probabilistic Taxonomy for Text Understanding
Wentao Wu, University of Wisconsin-Madison; Hongsong Li, Microsoft Research Asia; Haixun Wang, Microsoft Research Asia; Kenny Q. Zhu, Shanghai Jiao Tong University

Optimizing Index for Taxonomy Keyword Search
Bolin Ding, University of Illinois at Urbana Champaign; Haixun Wang, Microsoft Research Asia; Ruoming Jin, Kent State University; Jiawei Han, University of Illinois at Urbana Champaign; Zhongyuan Wang, Microsoft Research Asia

SIGMOD Research 15: Social Networks and Graph Databases II

Location: Arizona Ballroom V
Hashtag: #sigmod12 #R15Px
Session Chair: Wook-Shin Han (Kyungpook National University)

A Model-Based Approach to Attributed Graph Clustering
Zhiqiang Xu, Nanyang Technological University; Yiping Ke, Institute of High Performance Computing, Singapore; Yi Wang, National University of Singapore; Hong Cheng, The Chinese University of Hong Kong; James Cheng, Nanyang Technological University

Towards Effective Partition Management for Large Graphs
Shengqi Yang, University of California, Santa Barbara; Xifeng Yan, University of California, Santa Barbara; Bo Zong, University of California, Santa Barbara; Arijit Khan, University of California, Santa Barbara

TreeSpan: Efficiently Computing Similarity All-Matching
Gaoping Zhu, University of New South Wales; Xuemin Lin, University of New South Wales; Ke Zhu, University of New South Wales; Wenjie Zhang, University of New South Wales; Jeffrey Xu Yu, The Chinese University of Hong Kong

SIGMOD Research 16: Indexing and Physical Database Design II

Location: Arizona Ballroom VI–VIII

Hashtag: #sigmod12 #R16Px

Session Chair: Atish Das Sarma (Google Research)

Locality-Sensitive Hashing Scheme Based on Dynamic Collision Counting

Junho Gan, Sun Yat-Sen University; Jianlin Feng, Sun Yat-Sen University; Qiong Fang, Hong Kong University of Science and Technology; Wilfred Ng, Hong Kong University of Science and Technology

Efficient External-Memory Bisimulation on DAGs

Jelle Hellings, Hasselt University and Transnational University of Limburg; George H. L. Fletcher, Eindhoven University of Technology; Herman Haverkort, Eindhoven University of Technology

Materialized View Selection for XQuery Workloads

Asterios Katsifodimos, INRIA Saclay and Université Paris-Sud; Ioana Manolescu, INRIA Saclay and Université Paris-Sud; Vasilis Vassalos, Athens University of Economics and Business

SIGMOD Demonstrations A: Information Extraction, Search, Performance, and Clouds

Location: Vaquero Ballroom A

Hashtag: #sigmod12 #demoAx

Automatic Web-Scale Information Extraction

Philip Bohannon, Yahoo! Research; Nilesh Dalvi, Yahoo! Research; Yuval Filmus, University of Toronto; Nori Jacoby, Yahoo!; Sathiya Keerthi, Yahoo! Research; Alok Kirpal, Yahoo! Research

Just-in-Time Information Extraction using Extraction Views

Amr El-Helw, EMC Corp.; Mina Farid, University of Waterloo; Ihab Ilyas, Qatar Computing Research Institute

ColumbuScout: Towards Building Local Search Engines over Large Databases

Cody Hansen, University of Utah; Feifei Li, University of Utah

Sofia Search: A Tool for Automating Related-Work Search

Behzad Golshan, Boston University; Theodoros Lappas, Boston University; Evimaria Terzi, Boston University

RACE: Real-Time Applications over Cloud-Edge

Badrish Chandramouli, Microsoft Research; Joris Claessens, Microsoft Research; Suman Nath, Microsoft Research; Ivo Santos, Microsoft Research; Wenchao Zhou, University of Pennsylvania

Partiqle: An Elastic SQL Engine over Key-Value Stores

Junichi Tatemura, NEC Laboratories America; Oliver Po, NEC Laboratories America; Wang-Pin Hsiung, NEC Laboratories America; Hakan Hacigümüs, NEC Laboratories America

JustMyFriends: Full SQL, Full Transactional Amenities, and Access Privacy

Arthur Meacham, New York University; Dennis Shasha, New York University

Dynamic Optimization of Generalized SQL Queries with Horizontal Aggregations

Carlos Ordonez, University of Houston; Javier García-García, UNAM; Zhibo Chen, University of Houston

ConsAD: A Real-Time Consistency Anomalies Detector

Kamal Zellag, McGill University; Bettina Kemme, McGill University

Interactive Performance Monitoring of a Composite OLTP and OLAP Workload

Anja Bog, Hasso Plattner Institute, University of Potsdam; Kai Sachs, SAP AG; Hasso Plattner, Hasso Plattner Institute, University of Potsdam

SIGMOD Tutorial 7: Differential Privacy in Data Publication and Analysis (Part 2)

Location: Arizona Ballroom IV

Hashtag: #sigmod12 #tutorial7

Presenters: Yin Yang, Advanced Digital Sciences Center, Singapore; Zhenjie Zhang, Advanced Digital Sciences Center, Singapore; Gerome Miklau, University of Massachusetts, Amherst; Marianne Winslett, University of Illinois at Urbana-Champaign; Xiaokui Xiao, Nanyang Technological University

AWARDS

SIGMOD Contributions Award

For strengthening and humanizing the database community by originating and developing the "Distinguished Profile in Database Research" series.

Marianne's efforts to interview a large number of senior members as well as emerging stars of our community have been nothing short of heroic. These interviews bring home to us the story of their lives including the key decision points along the way, their reflections on the state of our community, and advice to the next generation of researchers. The transcripts and video tapes of these interviews will prove to be a treasure trove of insights for many future generations of young database researchers.

Marianne Winslett has been a professor in the Department of Computer Science at the University of Illinois since 1987. She is an ACM Fellow and the recipient of a Presidential Young Investigator Award from the US National Science Foundation. She is the former vice-chair of ACM SIGMOD and has served on the editorial boards of ACM Transactions on the Web, ACM Transactions on Database Systems, IEEE Transactions on Knowledge and Data Engineering, ACM Transactions on Information and Systems Security, and the Very Large Data Bases Journal. She has received two best paper awards for research on managing regulatory compliance data (VLDB, SSS), one best paper award for research on analyzing browser extensions to detect security vulnerabilities (Usenix Security), and one for keyword search (ICDE). Her PhD is from Stanford University.

SIGMOD Test-of-Time Award

Executing SQL over Encrypted Data in the Database-Service-Provider Model

Hakan Hacigumus, Bala Iyer, Chen Li, Sharad Mehrotra

Rapid advances in networking and Internet technologies have fueled the emergence of the "software as a service" model for enterprise computing. Successful examples of commercially viable software services include rent-a-spreadsheet, electronic mail services, general storage services, disaster protection services. "Database as a Service" model provides users power to create, store, modify, and retrieve data from anywhere in the world, as long as they have access to the Internet. It introduces several challenges, an important issue being data privacy. It is in this context that we specifically address the issue of data privacy. There are two main privacy issues. First, the owner of the data needs to be assured that the data stored on the service-provider site is protected against data thefts from outsiders. Second, data needs to be protected even from the service providers, if the providers themselves cannot be trusted. In this paper, we focus on the second challenge. Specifically, we explore techniques to execute SQL queries over encrypted data. Our strategy is to process as much of the query as possible at the service providers' site, without having to decrypt the data. Decryption and the remainder of the query processing are performed at the client site. The paper explores an algebraic framework to split the query to minimize the computation at the client site. Results of experiments validating our approach are also presented.

This paper from the SIGMOD 2002 Conference remarkably anticipated the world of "Database as Service" which did come about and continues to grow in importance. To get a sense of how visionary the work was, consider that this paper was published in June 2002 (and thus accepted in Jan 2002), even a couple of months before Amazon EC2 and S3 services were launched (of course, Amazon RDS

and SQL Azure came much later). The core of the paper focuses on the challenges of how to leverage cloud services while keeping some of the information (at the discretion of the enterprise/user) hidden from the service provider. Beyond the specific algorithmic details, the key contribution is the framework: (i) introduction of a mapping function, and (ii) query splitting logic to ensure how the work can be distributed across cloud and client when some information is encrypted. Is this framework used by enterprises today? As best as we can tell, the answer is perhaps no. But, is the framework interesting and has real possibilities of adoption and further impact and more follow-on by research community? Absolutely. In summary, this paper is one of the early papers to foresee the world of Database as Service (before any one of us were working on that problem). The specific technical focus was dealt with reasonable depth. The impact of the technical focus has not yet been seen by the industry but this paper has the possibility of inspiring much more follow-on work/thinking (beyond 140+ citations it already has in ACM Digital Library).

Hakan Hacigumus is the head of Data Management Research at NEC Labs America. His current interests include data management in the cloud, big data, data analytics, mobility, and service oriented business models. Prior to NEC Labs, he was a researcher at IBM Almaden Research Center, where worked on a wide range of areas in data management and services research. He received his Ph.D. in Computer Science from the University of California, Irvine.

Balakrishna (Bala) Iyer works for IBM as a Distinguished Engineer for Database Technology. He earned his B.Tech from IIT -Bombay, MS and PhD degrees from Rice University. He has worked previously for Bell Labs, Murray Hill, NJ. Bala has made contributions to the field of database in the area of temporal data, database as a service compression, sorting, query processing, data mining, encoded vector representation and processing. Many of his innovation are used every day, having been incorporated in IBM's data management products like VSAM, IMS, DB2 and IBM Intelligent Miner, and products from other leading vendors. His work on the temporal data model led to the standardization of temporal function in SQL 2011.

Chen Li is an associate professor in the Department of Computer Science at the University of California, Irvine. He received his Ph.D.degree in Computer Science from Stanford University in 2001, and his M.S. and B.S. in Computer Science from Tsinghua University, China, in 1996 and 1994, respectively. He received a National Science Foundation CAREER Award in 2003 and many other NSF grants and industry gifts. He was once a part-time Visiting Research Scientist at Google.

His research interests are in the fields of data management and information search, including text search, data-intensive computing, and data integration. He is the founder of Bimaple Technology Inc., a company providing powerful search for enterprises and developers.

Sharad Mehrotra is a Professor in the School of Information and Computer Science at University of California, Irvine and founding Director of the Center for Emergency Response Technologies (CERT) at UCI. From 2002-2009 he served as the Director and PI of the RESCUE project (Responding to Crisis and Unexpected Events) which, funded by NSF through its large ITR program, spanned 7 schools and consisted of 60 members. He is the recipient of Outstanding Graduate Student Mentor Award in 2005. Prior to joining UCI, he was a member of the faculty at University of Illinois, Urbana Champaign in the Department of Computer Science where he was the recipient of the C. W. Gear Outstanding Junior Faculty Award. Mehrotra has also served as a Scientist at Matsushita Information Technology Laboratory immediately after graduating with a Ph.D. from University of Texas at Austin (1988-1993).

Mehrotra's research expertise is in data management and distributed systems areas in which he has made many pioneering contributions. Two such contributions include the concept of "database as a service" and "use of information retrieval

techniques, particularly relevance feedback, in multimedia search". Mehrotra is a recipient of numerous best paper nominations and awards including SIGMOD Best Paper award in 2001 for a paper entitled "Locally Adaptive Dimensionality Reduction for Indexing Large Time Series Databases", and best paper award in DASFAA 2004 for the paper entitled "Efficient Execution of Aggregation Queries over Encrypted Databases". Another of his paper entitled "Concurrency Control in Hierarchical Multidatabase System" was selected as best of VLDB 1994 submissions invited for the VLDB Journal. Mehrotra's recent research focuses on data quality, data privacy particularly in the context of cloud computing and sensor driven situational awareness systems.

SIGMOD Best Paper Award

High-Performance Complex Event Processing over XML Streams

Barzan Mozafari, University of California, Los Angeles; Kai Zeng, University of California, Los Angeles; Carlo Zaniolo, University of California, Los Angeles

Much research attention has been given to delivering high-performance systems that are capable of complex event processing (CEP) in a wide range of applications. However, many current CEP systems focus on processing efficiently data having a simple structure, and are otherwise limited in their ability to support efficiently complex continuous queries on structured or semi-structured information. However, XML streams represent a very popular form of data exchange, comprising large portions of social network and RSS feeds, financial records, configuration files, and similar applications requiring advanced CEP queries. In this paper, we present the XSeq language and system that support CEP on XML streams, via an extension of XPath that is both powerful and amenable to an efficient implementation. Specifically, the XSeq language extends XPath with natural operators to express sequential and Kleene-* patterns over XML streams, while remaining highly amenable to efficient implementation. XSeq is designed to take full advantage of recent advances in the field of automata on Visibly Pushdown Automata (VPA), where higher expressive power can be achieved without compromising efficiency (whereas the amenability to efficient implementation was not demonstrated in XPath extensions previously proposed).
We illustrate XSeq's power for CEP applications through examples from different domains, and provide formal results on its expressiveness and complexity. Finally, we present several optimization techniques for XSeq queries. Our extensive experiments indicate that XSeq brings outstanding performance to CEP applications: two orders of magnitude improvement are obtained over the same queries executed in general-purpose XML engines.

Barzan Mozafari is currently a Postdoc Associate at Massachusetts Institute of Technology. He earned his PhD in Computer Science from the University of California at Los Angeles, where he worked on scalable solutions for pattern discovery and detection from large volumes of data, meeting several system, language and algorithmic challenges. His research interests include distributed databases, machine learning, crowd-sourcing and cloud computing.

Kai Zeng received the bachelor's degree in computer science from Zhejiang University, China, in 2009. He is currently working toward the PhD degree in database systems, under the supervision of Professor Carlo Zaniolo. He is also a research assistant. His research interests include query processing, pattern matching in data streams and massive data.

Carlo Zaniolo is a professor of Computer Science at UCLA where he occupies the N.E. Friedmann chair in Knowledge Science. His research interests include Data Stream Management Systems, Data Mining, Logic Based Languages, and Web Information Systems.

Edgar F. Codd Innovations Award

For innovative and highly significant contributions of enduring value to the development, understanding, or use of database systems and databases.

Bruce Lindsay *has been a leader and inventor in many of the key systems initiatives in the data management field. As a member of the original System R team, the R* project, the Starburst project, and then several content management projects, Bruce has created fundamental technologies in a broad set of database areas, including core relational databases (authorization, high performance transactions, locking and deadlock detection), extensible databases (object management, type management, production rules for query processing), distributed databases (snapshots, distributed DDL, presumed commit, presumed abort, distributed query processing), and management of unstructured data (XML, novel indexing). He thinks broadly and has uncanny intuition for the system-level issues that has led to his innovations to have lasting impact on commercial database products.*

SIGMOD Jim Gray Doctoral Dissertation Award

ACM SIGMOD is pleased to present the 2012 SIGMOD Jim Gray Doctoral Dissertation Award to F. Ryan Johnson. Johnson completed his dissertation titled "Scalable Storage Managers for the Multicore Era" at Carnegie Mellon University. Johnson's dissertation is a tour de force in identifying bottlenecks when scaling OLTP systems to many cores, proposing innovative solutions to each of them. The ideas in the thesis such as speculative lock inheritance, new techniques for combining log requests, and data-oriented transaction execution are highly innovative, and the work is remarkable for its breadth, depth, thorough implementation, and evaluation.

Ryan Johnson is an Assistant Professor at the University of Toronto specializing in systems aspects of database engines, particularly in the context of modern hardware. He graduated with M.S. and PhD degrees in Computer Engineering from Carnegie Mellon University in 2010, after completing a B.S. in Computer Engineering at Brigham Young University in 2004. In addition to his work with database systems, Johnson has interests in computer architecture, operating systems, compilers, and hardware design.

SIGMOD Jim Gray Doctoral Dissertation Honorable Mention

ACM SIGMOD is also pleased to recognize Bogdan Alexe for an Honorable Mention for the 2012 SIGMOD Jim Gray Doctoral Dissertation Award. Alexe completed his dissertation titled "Interactive and Modular Design of Schema Mappings" at the University of California, Santa Cruz. Alexe's dissertation makes substantial contributions to the important problem of designing schema mappings through novel principled algorithms and the first benchmark in this area.

Bogdan Alexe is a researcher at IBM Research - Almaden. His work focuses on large scale entity resolution and integration. His past research covered topics in information integration, data exchange and schema mappings. Bogdan graduated with a Ph.D. from University of California at Santa Cruz, and an M.Sc. from Ecole Polytechnique/Telecom ParisTech, both in Computer Science.

PODS Best Paper Award

Best Paper Award: Worst-Case Optimal Join Algorithms

Hung Q. Ngo, University at Buffalo, SUNY; Ely Porat, Bar-Ilan University; Christopher Ré, University of Wisconsin-Madison; Atri Rudra, University at Buffalo, SUNY

Efficient join processing is one of the most fundamental and well-studied tasks in database research. In this work, we examine algorithms for natural join queries over many relations and describe a novel algorithm to process these queries optimally in terms of worst-case data complexity. Our result builds on recent work by Atserias, Grohe, and Marx, who gave bounds on the size of a full conjunctive query in terms of the sizes of the individual relations in the body of the query. These bounds, however, are not constructive: they rely on Shearer's entropy inequality which is information-theoretic. Thus, the previous results leave open the question of whether there exist algorithms whose running time achieve these optimal bounds. An answer to this question may be interesting to database practice, as we show in this paper that any project-join plan is polynomially slower than the optimal bound for some queries. We construct an algorithm whose running time is worst-case optimal for all natural join queries. Our result may be of independent interest, as our algorithm also yields a constructive proof of the general fractional cover bound by Atserias, Grohe, and Marx without using Shearer's inequality. In addition, we show that this bound is equivalent to a geometric inequality by Bollobs and Thomason, one of whose special cases is the famous Loomis-Whitney inequality. Hence, our results algorithmically prove these inequalities as well. Finally, we discuss how our algorithm can be used to compute a relaxed notion of joins.

Hung Q. Ngo is an Associate Professor at the Computer Science and Engineering department, State University of New York (SUNY) at Buffalo. He received a Ph.D. in Computer Science and an M.S. in Mathematics from the University of Minnesota, Twin Cities. His main research interests are in the theory of switching networks and algorithmic group testing.

Ely Porat is an Associate Professor at Bar-Ilan University. He received his Doctorate at Bar-Ilan University in 2000. Following that, he fulfilled his military service and, in parallel, worked as a faculty member at Bar-Ilan University. Porat spent the spring 2007 semester as a Visiting Scientist in Google Mountain View. He is a consultor to Google in Tel Aviv, and holds positions as a visiting professor at the University of Michigan and at Tel Aviv University.

Christopher (Chris) Ré is an assistant professor in the department of Computer Sciences at the University of Wisconsin-Madison. The goal of his work is to enable users and developers to build applications that more deeply understand and exploit data. Chris received his PhD from the University of Washington, Seattle under the supervision of Dan Suciu. For his PhD work in the area of probabilistic data management, Chris received the SIGMOD 2010 Jim Gray Dissertation Award. Chris received an NSF CAREER Award in 2011.

Atri Rudra is an Assistant Professor of Computer Science and Engineering at University at Buffalo, State University of New York, Buffalo. Atri received his Bachelor's degree from Indian Institute of Technology, Kharagpur, India in 2000 and his Ph.D. from University of Washington in 2007. From 2000-2002, he was a Research Staff Member at IBM India Research Lab, New Delhi, India. His research interests lie in theoretical computer science and in particular, theory of error-correcting codes, data stream and sub-linear algorithms, game theory and algorithmic mechanism design, approximation algorithms, computational complexity, finite field theory and applications. He is a recipient of an NSF CAREER award (2009), HP Labs Innovation Research Award (2010), ESA best paper award (2010) and the UB Exceptional Scholars - Young Investigator award (2011).

ACM PODS Alberto O. Mendelzon Test-of-Time Award 2012

Containment and Equivalence for an XPath Fragment

Gerome Miklau, University of Massachusetts, Amherst; Dan Suciu, University of Washington

XPath is a simple language for navigating an XML document and selecting a set of element nodes. XPath expressions are used to query XML data, describe key constraints, express transformations, and reference elements in remote documents. This paper studies the containment and equivalence problems for a fragment of the XPath query language, with applications in all these contexts.

In particular, we study a class of XPath queries that contain branching, label wildcards and can express descendant relationships between nodes. Prior work has shown that languages which combine any two of these three features have efficient containment algorithms. However, we show that for the combination of features, containment is coNP-complete. We provide a sound and complete EXPTIME algorithm for containment, and study parameterized PTIME special cases. While we identify two parameterized classes of queries for which containment can be decided efficiently, we also show that even with some bounded parameters, containment is coNP-complete. In response to these negative results, we describe a sound algorithm which is efficient for all queries, but may return false negatives in some cases.

"The paper studied static analysis problems for XPath, a query language at the core of processing XML documents and XML document databases. The results of Miklau and Suciu shed light on the boundary between tractability and intractability for XPath query containment, since it was previously known that the containment problem was solvable in polynomial time for XPath queries in which any two of the three navigation axes of XPath are allowed. Both the paper in the PODS 2002 proceedings and its subsequent full version in the Journal of the Association for Computing Machinery have received hundreds of citations each. Moreover, this work initiated a fruitful line of research on the static analysis of XML query languages that brought together researchers from database theory and automata theory."

Gerome Miklau is an Associate Professor at the University of Massachusetts, Amherst. His primary research interest is the secure management of large-scale data. This includes evaluating threats to privacy in published data, devising techniques for the safe publication of social networks, network traces, and audit logs, as well as designing database management systems to implement security policies. He was awarded a Lilly Teaching Fellowship in 2011, an NSF CAREER Award in 2007, and he won the 2006 ACM SIGMOD Dissertation Award. He received his Ph.D. in Computer Science from the University of Washington in 2005. He earned Bachelor's degrees in Mathematics and in Rhetoric from the University of California, Berkeley, in 1995.

Dan Suciu is a Professor in Computer Science at the University of Washington. He received his Ph.D. from the University of Pennsylvania in 1995, was a principal member of the technical staff at AT&T Labs and joined the University of Washington in 2000. Suciu is conducting research in data management, with an emphasis on topics related to Big Data and data sharing, such as probabilistic data, data pricing, parallel data processing, data security. He is a co-author of two books Data on the Web: from Relations to Semistructured Data and XML, 1999, and Probabilistic Databases, 2011. He is a Fellow of the ACM, holds twelve US patents, received the 2000 ACM SIGMOD Best Paper Award, the 2010 PODS Ten Years Best paper award, and is a recipient of the NSF Career Award and of an Alfred P. Sloan Fellowship. Suciu serves on the VLDB Board of Trustees, and is an associate editor for the VLDB Journal, for ACM TOIS, and for Information Systems. Suciu's PhD students Gerome Miklau and Christopher Re received the ACM SIGMOD Best Dissertation Award in 2006 and 2010 respectively, and Nilesh Dalvi was a runner up in 2008.

KEYNOTES

PODS Keynote: What Next? A Half-Dozen Data Management Research Goals for Big Data and the Cloud

Surajit Chaudhuri (Microsoft Research)

"Big Data" and the Cloud are two disruptions that are influencing our field today. In this talk, I will outline the nature of this disruption. Next, following the structure of Jim Gray's Turing award lecture, I will describe six fundamental technical challenges that will be important as a research community to address in order to take advantage of these disruptions. While some of the challenges are unique to these disruptions, others are known challenges but whose importance is amplified by Big Data and the Cloud. A good solution to several of these problems will require a close interaction between data management systems and theory sub-communities.

Surajit Chaudhuri is a Distinguished Scientist at Microsoft research. His current areas of interest are enterprise data analytics, self-manageability and multi-tenant technology for cloud database services. 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. Surajit is an ACM Fellow, a recipient of the ACM SIGMOD Edgar F. Codd Innovations Award, ACM SIGMOD Contributions Award, a VLDB 10 year Best Paper Award, and an IEEE Data Engineering Influential Paper Award. He was the Program Committee Chair for ACM SIGMOD 2006, 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.

SIGMOD Keynote Talk 1: Analytic Database Technologies for a New Kind of User - The Data Enthusiast

Pat Hanrahan, Stanford University and Tableau Software

Analytics enables businesses to increase the efficiency of their activities and ultimately increase their profitability. As a result, it is one of the fastest growing segments of the database industry. There are two usages of the word analytics. The first refers to a set of algorithms and technologies, inspired by data mining, computational statistics, and machine learning, for supporting statistical inference and prediction. The second is equally important: analytical thinking. Analytical thinking is a structured approach to reasoning and decision making based on facts and data. Most of the recent work in the database community has focused on the first, the algorithmic and systems problems. The people behind these advances comprise a new generation of data scientists who have either the mathematical skills to develop advanced statistical models, or the computer skills to develop or implement scalable systems for processing large, complex datasets. The second aspect of analytics -- supporting the analytical thinker -- although equally important and challenging, has received much less attention.

In this talk, I will describe recent advances in in making both forms of analytics accessible to a broader range of people, who I call data enthusiasts. A data enthusiast is an educated person who believes that data can be used to answer a question or solve a problem. These people are not mathematicians or programmers, and only know a bit of statistics. I'll review recent work on building easy-to-use, yet powerful, visual interfaces for working with data; and the analytical database technology needed to support these interfaces.

Pat Hanrahan is the CANON Professor of Computer Science and Electrical Engineering at Stanford University where he teaches computer graphics.His current research involves visualization, image synthesis, and graphics systems and architectures. Before joining Stanford he was a faculty member at Princeton. Pat has also worked at Pixar where he developed developed volume rendering software and was the chief architect of the RenderMan(TM) Interface - a protocol that allows modeling programs to describe scenes to high quality rendering programs. In addition to PIXAR, he has founded two companies, Tableau and PeakStream, and served on the technical advisory boards of NVIDIA, Exluna, Neoptica, VSee, Procedural and Skytree. Professor Hanrahan has received three university teaching awards.He has received two Academy Awards for Science and Technology, the Spirit of America Creativity Award, the SIGGRAPH Computer Graphics Achievement Award, the SIGGRAPH Stephen A. Coons Award, and the IEEE Visualization Career Award. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences.

SIGMOD Keynote Talk 2: Symbiosis in Scale Out Networking and Data Management

Amin Vahdat (University of California San Diego and Google)

This talk highlights the symbiotic relationship between data management and networking through a study of two seemingly independent trends in the traditionally separate communities: large-scale data processing and software defined networking. First, data processing at scale increasingly runs across hundreds or thousands of servers. We show that balancing network performance with computation and storage is a prerequisite to both efficient and scalable data processing. We illustrate the need for scale out networking in support of data management through a case study of TritonSort, currently the record holder for several sorting benchmarks, including GraySort and JouleSort. Our TritonSort experience shows that disk-bound workloads require 10 Gb/s provisioned bandwidth to keep up with modern processors while emerging flash workloads require 40 Gb/s fabrics at scale.

We next argue for the need to apply data management techniques to enable *Software Defined Networking* (SDN) and *Scale Out Networking*. SDN promises the abstraction of a single logical network fabric rather than a collection of thousands of individual boxes. In turn, scale out networking allows network capacity (ports, bandwidth) to be expanded incrementally, rather than by wholesale fabric replacement. However, SDN requires an extensible model of both static and dynamic network properties and the ability to deliver dynamic updates to a range of network applications in a fault tolerant and low latency manner. Doing so in networking environments where updates are typically performed by timer-based broadcasts and models are specified as comma-separated text files processed by one-off scripts presents interesting challenges. For example, consider an environment where applications from routing to traffic engineering to monitoring to intrusion/anomaly detection all essentially boil down to inserting, triggering and retrieving updates to/from a shared, extensible data store.

Amin Vahdat is a Principal Engineer at Google working on data center and wide-area network architecture. He is also a Professor and holds the Science Applications International Corporation Chair in the Department of Computer Science and Engineering at the University of California San Diego. Vahdat's research focuses broadly on computer systems, including distributed systems, networks, and operating systems. He received a PhD in Computer Science from UC Berkeley under the supervision of Thomas Anderson after spending the last year and a half as a Research Associate at the University of Washington. Vahdat is an ACM Fellow and a past recipient of the the NSF CAREER award, the Alfred P. Sloan Fellowship, and the Duke University David and Janet Vaughn Teaching Award.

TUTORIALS

PODS Tutorial 1: Linguistic Foundations for Bidirectional Transformations
Presenter: Benjamin C. Pierce, University of Pennsylvania

Computing is full of situations where two different structures must be "connected" in such a way that updates to each can be propagated to the other. This is a generalization of the classical view update problem, which has been studied for decades in the database community; more recently, related problems have attracted considerable interest in other areas, including programming languages, software model transformation, user interfaces, and system configuration. Among the fruits of this cross-pollination has been the development of a linguistic perspective on the problem. Rather than taking some view definition language as fixed (e.g., choosing some subset of relational algebra) and looking for tractable ways of "inverting" view definitions to propagate updates from view to source, we can directly design new bidirectional programming languages in which every expression denotes a pair of functions mapping updates on one structure to updates on the other. Such structures are often called lenses. The foundational theory of lenses has been studied extensively, and lens-based language designs have been developed in several domains, including strings, trees, relations, graphs, and software models. These languages share some common elements with modern functional languages – in particular, they come with very expressive type systems. In other respects, they are rather novel and surprising. This tutorial surveys recent developments in the theory of lenses and the practice of bidirectional programming languages.

Benjamin Pierce joined the CIS Department at Penn in 1998. Previously, he was on the faculty at Indiana University and held research fellowships at Cambridge University, the University of Edinburgh, and INRIA-Roquencourt. He received his Ph.D. in Computer Science at Carnegie Mellon University in 1991. His research centers on programming languages, static type systems, concurrent and distributed programming, and synchronization technologies. His books include the widely used graduate text Types and Programming Languages. *He is also the lead designer of the popular* Unison *file synchronizer.*

PODS Tutorial 2: Approximate Computation and Implicit Regularization for Very Large-scale Data Analysis
Presenter: Michael W. Mahoney, Stanford University

Database theory and database practice are typically done by computer scientists who adopt what may be termed an algorithmic perspective on their data. This perspective is very different than the perspective adopted by statisticians, scientific computers, machine learners, and other who work on what may be broadly termed statistical data analysis. I will address fundamental aspects of this algorithmic-statistical disconnect, with an eye to bridging the gap between these two very different approaches. A concept that lies at the heart of this disconnect is that of statistical regularization, a notion that has to do with how robust is the output of an algorithm to the noise properties of the input data. Although it is nearly completely absent from computer science, which historically has taken the input data as given and modeled algorithms discretely, regularization in one form or another is central to nearly every application domain that applies algorithms to noisy data. By using several case studies, I will illustrate, both theoretically and empirically, the nonobvious fact that approximate computation, in and of itself, can implicitly lead to statistical regularization. This and other recent work suggests that, by exploiting in a more principled way the statistical properties implicit in worst-case algorithms, one can in many cases satisfy the bicriteria of having algorithms that are scalable to very large-scale databases and that also have good inferential or predictive properties.

Michael Mahoney is at Stanford University. His research interests center around algorithms for very large-scale statistical data analysis, including both theoretical and applied aspects of problems in scientific and Internet domains. His current research interests include geometric network analysis; developing approximate computation and regularization methods for large informatics graphs; applications to community detection, clustering, and information dynamics in large social and information networks; and the theory of randomized matrix algorithms and its application to genetics, medical imaging, and Internet problems. He has been a faculty member at Yale University and a researcher at Yahoo, and his PhD was is computational statistical mechanics at Yale University.

SIGMOD Tutorial 1: Mob Data Sourcing

Presenters: Daniel Deutch, Ben Gurion University; Tova Milo, Tel Aviv University

Crowdsourcing is an emerging paradigm that harnesses a mass of users to perform various types of tasks. We focus in this tutorial on a particular form of crowdsourcing, namely crowd (or mob) datasourcing whose goal is to obtain, aggregate or process data. We overview crowd datasourcing solutions in various contexts, explain the need for a principled solution, describe advances towards achieving such a solution, and highlight remaining gaps.

Daniel Deutch is an Assistant Professor in the Computer Science Department of Ben Gurion University. He has received his PhD degree in Computer Science from Tel Aviv University in 2010 and was a Postdoc at the University of Pennsylvania (UPenn) and the INRIA research institute. His research interests focus on web data management, in particular using probabilistic models and techniques. During his PhD studies Daniel has received a number of awards for his research, including the Israeli Ministry of Science Eshkol grant and ICDT best student paper award. Daniel has been a member of the program committee of various international conferences and workshops (including WWW, ICDT, PODS). He has received research grants from the US-Israel Binational Science Foundation and the Israeli Ministry of Science.

Tova Milo received her Ph.D. degree in Computer Science from the Hebrew University, Jerusalem, in 1992. After graduating she worked at the INRIA research institute in Paris and at University of Toronto and returned to Israel in 1995, joining the School of Computer Science at Tel Aviv University where she is now a full Professor and Department head. Her research focuses on advanced database applications such 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 Chair of 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 the European Union. She is a recipient of the 2010 ACM PODS Alberto O. Mendelzon Test-of-Time Award and of the prestigious EU ERC Advanced Investigators grant.

SIGMOD Tutorial 2: Managing and Mining Large Graphs: Patterns and Algorithms

Presenters: Christos Faloutsos, Carnegie Mellon University; U Kang, Carnegie Mellon University

Graphs are everywhere: social networks, the World Wide Web, biological networks, and many more. The sizes of graphs are growing at unprecedented rate, spanning millions and billions of nodes and edges. What are the patterns in large graphs, spanning Giga, Tera, and heading toward Peta bytes? What are the best tools, and how can they help us solve graph mining problems? How do we scale up algorithms for handling graphs with billions of nodes and edges? These are exactly the goals of this tutorial. We start with the patterns in real-world static, weighted, and dynamic graphs. Then we describe important tools for large graph mining, including singular value decomposition, and Hadoop. Finally, we conclude with the design and the implementation of scalable graph mining algorithms on Hadoop. This tutorial is complementary to the related tutorial "Managing and Mining Large Graphs: Systems and Implementations".

Christos Faloutsos is a Professor at Carnegie Mellon University.He has received the Presidential Young Investigator Award by the National Science Foundation (1989), the Research Contributions Award in ICDM 2006, the SIGKDD Innovations Award (2010), eighteen ``best paper'' awards (including two ``test of time'' awards),and four teaching awards. He is an ACM Fellow,he has served as a member of the executive committee of SIGKDD;he has published over 200 refereed articles, 11 book chapters and one monograph. He holds six patents and he has given over 30 tutorials and over 10 invited distinguished lectures.His research interests include data mining for graphs and streams, fractals, database performance, and indexing for multimedia and bio-informatics data.

U Kang is a Ph.D. candidate in Computer Science at Carnegie Mellon University.He received B.S. in Computer Science and Engineering at Seoul National University.He won two best paper awards.He has published 16 refereed articles in major data mining and database venues. He holds two U.S. patents. His research interests include data mining in massive graphs.

SIGMOD Tutorial 3: Managing and Mining Large Graphs: Systems and Implementations

Presenters: Bin Shao, Microsoft Research Asia; Haixun Wang, Microsoft Research Asia; Yanhua Xiao, Fudan University and Microsoft Research Asia

We are facing challenges at all levels ranging from infrastructures to programming models for managing and mining large graphs. A lot of algorithms on graphs are ad-hoc in the sense that each of them assumes that the underlying graph data can be organized in a certain way that maximizes the performance of the algorithm. In other words, there is no standard graph systems based on which graph algorithms are developed and optimized. In response to this situation, a lot of graph systems have been proposed recently. In this tutorial, we discuss several representative systems. Still, we focus on providing perspectives from a variety of standpoints on the goals and the means for developing a general purpose graph system. We highlight the challenges posed by the graph data, the constraints of architectural design, the different types of application needs, and the power of different programming models that support such needs.This tutorial is complementary to the related tutorial ``Managing and Mining Large Graphs: Patterns and Algorithms".

Bin Shao is an associate researcher at Microsoft Research Asia. He received his B.E. (2005) in Computer Science from Shandong University and Ph.D. in Computer Science (2010) from Fudan University. His Ph.D. research is operational transformation (OT), CSCW, and optimistic consistency maintenance. His work enables OT techniques to be applied in the next generation of mobile and Web based collaborative applications. He has been working on a distributed in-memory graph engine called Trinity after joining Microsoft Research Asia. His research interests include graph database, scalable distributed systems, and all-in-memory system design and performance optimization.

Haixun Wang is a senior researcher at Microsoft Research Asia in Beijing, China, where he manages the group of Data Management, Analytics, and Services. Before joining Microsoft, he had been a research staff member at IBM T. J. Watson Research Center for 9 years. He was Technical Assistant to Stuart Feldman (Vice President of Computer Science of IBM Research) from 2006 to 2007, and Technical Assistant to Mark Wegman (Head of Computer Science of IBM Research) from 2007 to 2009. Haixun Wang has published more than 120 research papers in referred international journals and conference proceedings. He is on the editorial board of Distributed and Parallel Databases (DAPD), IEEE Transactions of Knowledge and Data Engineering (TKDE), Knowledge and Information System (KAIS), Journal of Computer Science and Technology (JCST). He is PC co-Chair of CIKM 2012, ICMLA 2011, WAIM 2011. Haixun Wang got the ER 2008 Conference best paper award (DKE 25 year award), and ICDM 2009 Best Student Paper run-up award.

Yanghua Xiao is an assistant professor of computer science at Fudan University, China. He had been a visiting professor of Human Genome Sequencing Center at Baylor College of Medicine, and visiting researcher of Microsoft Research Asia. His research interests include graph database and graph data mining. In recent years, he has published more than 20 papers about graph data management and complex network analysis in international leading journals and top conferences, including SIGMOD, ICSE, ICDE, EDBT, Physical Review E and Pattern Recognition.

SIGMOD Tutorial 4: Computational Reproducibility: State-of-the-Art, Challenges, and Database Research Opportunities

Presenters: Juliana Freire, NYU-Poly; Philippe Bonnet, IT University of Copenhagen; Dennis Shasha, NYU

Computational experiments have become an integral part of the scientific method, but reproducing, archiving, and querying them is still a challenge. The first barrier to a wider adoption is the fact that it is hard both for authors to derive a compendium that encapsulates all the components needed to reproduce a result and for reviewers to verify the results. In this tutorial, we will present a series of guidelines and, through hands-on examples, review existing tools to help authors create of reproducible results. We will also outline open problems and new directions for database-related research having to do with querying computational experiments.

Juliana Freire is a Professor of Computer Science at NYU Poly. Her research interests include Web mining and crawling, large-scale information integration, information visualization, and scientific data management. She is a co-creator of VisTrails (www.vistrails.org), an open-source data analysis and visualization system that supports the creation and publication of reproducible results. Since 2010, she has been working with the repeatability initiative of SIGMOD.

Philippe Bonnet is an associate professor at IT University of Copenhagen. He is an experimental computer scientist; his research interests include flash-based database systems, sensor data management and computational repeatability. Philippe currently serves as chair for the SIGMOD and VLDB reproducibility committees.

Dennis Shasha is a professor of computer science at New York University where he works with biologists on pattern discovery for network inference; with physicists and financial people on algorithms for time series; on database applications in untrusted environments; on database tuning;and on computational reproducibility. He has been working with the repeatability and workability initiative of SIGMOD since 2008.

SIGMOD Tutorial 5: Database Techniques for Linked Data Management

Presenters: Andreas Harth, Karlsruhe Institute of Technology (KIT); Katja Hose, Max-Planck Institute for Informatics; Ralf Schenkel, Saarland University

Linked Data refers to data published in accordance with a number of principles rooted in web standards. In the past few years we have witnessed a tremendous growth in Linked Data publishing on the web, leading to tens of billions of data items published online. Querying the data is a key functionality required to make use of the wealth of rich interlinked data. The goal of the tutorial is to introduce, motivate, and detail techniques for querying heterogeneous structured data from across the web. Our tutorial aims to introduce database researchers and practitioners to the new publishing paradigm on the web, and show how the abundance of data published as Linked Data can serve as fertile ground for database research and experimentation. As such, the tutorial focuses on applying database techniques to processing Linked Data, such as optimized indexing and query processing methods in the centralized setting as well as distributed approaches for querying. At the same time, we make the connection from Linked Data best practices to established technologies in distributed databases and the concept of Dataspaces and show differences as well as commonalities between the fields.

Andreas Harth is a post-doctoral researcher at Institute AIFB at the Karlsruhe Institute of Technology. His research interests are large-scale data interoperation on the Semantic Web, Linked Data, knowledge representation, computational logic and user interaction on web data. Andreas has published over a dozen papers in these areas, and is author of several open source software systems. Two of his systems were awarded prizes at the Semantic Web Challenge co-located with the International Semantic Web Conference. Andreas was awarded his Ph.D. by the Digital Enterprise Research Institute (DERI) at the National University of Ireland, Galway. He holds a Dipl.-Inf. (FH) (a nearby equivalent to the U.S.-style M.Sc.) from Fachhochschule Würzburg. Andreas worked as intern at Fraunhofer Gesellschaft in Würzburg and at IBM's Silicon Valley Lab in San Jose, CA. He visited USC's Information Sciences Institute in Marina del Rey, CA as a research assistant. Andreas has participated in numerous EU and national projects, participated in various program committees, and has served in the W3C Semantic Web Best Practices and Deployment and Rules Interchange Format working groups. In addition, he served as program committee member of numerous conferences and is one of the co-organizers of the Consuming Linked Data (COLD) workshop series and of the Semantic Web Challenge.

Katja Hose is a post-doctoral researcher at the Max-Planck Institute for Informatics in Saarbrücken, Germany. She obtained a diploma (M.Sc.) in Computer Science from Ilmenau University of Technology, joined the Databases & Information Systems Group at Ilmenau University of Technology as a research associate, and received her doctoral degree in Computer Science in 2009. Afterwards, joined the Max-Planck Institute for Informatics in Saarbrücken. Her current research interests range from query processing and optimization in distributed systems, heterogeneous databases, and rank-aware query operators to Linked Data processing, information retrieval, and knowledge extraction.

Ralf Schenkel is senior researcher at the Max-Planck Institute for Informatics in Saarbrücken, Germany, and a research group leader at Saarland University. His research interests include efficient and effective search on structured, semistructured, and unstructured data; of particular interest are social networks and distributed knowledge sources, as well as large-scale, long-term web archiving.Ralf serves as co-chair of INEX, the Initiative for the Evaluation of XML Retrieval; co-organized the 3rd ESAIR workshop on exploiting semantic annotations for IR at CIKM 2010; and has served on many program committees in DB and IR, including SIGIR, WSDM, WWW, CIKM, ICDE, SIGMOD, and VLDB.

SIGMOD Tutorial 6-7: Differential Privacy in Data Publication and Analysis

Presenters: Yin Yang, Advanced Digital Sciences Center, Singapore; Zhenjie Zhang, Advanced Digital Sciences Center, Singapore; Gerome Miklau, University of Massachusetts, Amherst; Marianne Winslett, University of Illinois at Urbana-Champaign; Xiaokui Xiao, Nanyang Technological University

Data privacy has been an important research topic in the security, theory and database communities in the last few decades. However, many existing studies have restrictive assumptions regarding the adversary's prior knowledge, meaning that they preserve individuals' privacy only when the adversary has rather limited background information about the sensitive data, or only uses certain kinds of attacks. Recently, differential privacy has emerged as a new paradigm for privacy protection with very conservative assumptions about the adversary's prior knowledge. Since its proposal, differential privacy had been gaining attention in many fields of computer science, and is considered among the most promising paradigms for privacy-preserving data publication and analysis. In this tutorial, we will motivate its introduction as a replacement for other paradigms, present the basics of the differential privacy model from a database perspective, describe the state of the art in differential privacy research, explain the limitations and shortcomings of differential privacy, and discuss open problems for future research.

Yin "David" Yang is a research scientist at ADSC. His research interests lie in database security and query optimization. He has published several papers in renowned venues about query authentication in outsourced databases. In addition, he has designed efficient query processing methods in various contexts, including data streams, relational keyword search, spatial databases, web portals, and wireless sensor networks. Currently, David is working on a project related to differentially private databases, led by Professor Marianne Winslett.

Zhenjie Zhang is currently research scientist in Advanced Digital Sciences Center, Illinois at Singapore Pte. He received his Ph.D. in computer science from the School of Computing, National University of Singapore, in 2010. Before that, he graduated with a B.S. degree from the Department of Computer Science and Engineering, Fudan University, in 2004. He was visiting student at the Hong Kong University of Science and Technology in 2008 and a visiting student at AT&T Shannon Lab in 2009. Before joining the Advanced Digital Sciences Center in October 2010, he worked as a Research Assistant and Research Fellow at the National University of Singapore from 2008 to 2010. His research interests cover a wide spectrum in computer science, including real-time analytics, non-metric indexing, game theory and data privacy. He has served as a Program Committee member for VLDB 2012, ICDE 2012, WWW 2010, VLDB 2010, KDD 2010 and other conferences. He was the recipient of President's Graduate Fellowship of National University of Singapore in 2007.

Gerome Miklau is an Associate Professor at the University of Massachusetts, Amherst. His primary research interest is the secure management of large-scale data. This includes evaluating threats to privacy in published data, devising techniques for the safe publication of social networks, network traces, and audit logs, as well as designing database management systems to implement security policies. He was awarded a Lilly Teaching Fellowship in 2011, an NSF CAREER Award in 2007, and he won the 2006 ACM SIGMOD Dissertation Award. He received his Ph.D. in Computer Science from the University of Washington in 2005. He earned Bachelor's degrees in Mathematics and in Rhetoric from the University of California, Berkeley, in 1995.

Marianne Winslett has been a professor in the Department of Computer Science at the University of Illinois since 1987. She is an ACM Fellow and the recipient of a Presidential Young Investigator Award from the US National Science Foundation. She is the former vice-chair of ACM SIGMOD and has served on the editorial boards of ACM Transactions on the Web, ACM Transactions on Database Systems, IEEE Transactions on Knowledge and Data Engineering, ACM Transactions on Information and Systems Security, and the Very Large Data Bases Journal. She has received two best paper awards for research on managing regulatory compliance data (VLDB, SSS), one best paper award for research on analyzing browser extensions to detect security vulnerabilities (Usenix Security), and one for keyword search (ICDE). Her PhD is from Stanford University.

Xiaokui Xiao is a Nanyang Assistant Professor (NAP) at the School of Computer Engineering, Nanyang Technological University (NTU). He obtained a PhD in Computer Science from the Chinese University of Hong Kong in 2008. Before joining NTU in 2009, he was a postdoctoral associate at the Cornell University. He is part of the DANTE research group at NTU.

SIGMOD RESEARCH PAPER ABSTRACTS

SIGMOD Research 1: Distributed and Parallel Databases

Calvin: Fast Distributed Transactions for Partitioned Database Systems

Alexander Thomson, Yale University; Thaddeus Diamond, Yale University; Shu-Chun Weng, Yale University; Kun Ren, Yale University; Philip Shao, Yale University; Daniel J. Abadi, Yale University

Many distributed storage systems achieve high data access throughput via partitioning and replication, each system with its own advantages and tradeoffs. In order to achieve high scalability, however, today's systems generally reduce transactional support, disallowing single transactions from spanning multiple partitions. Calvin is a practical transaction scheduling and data replication layer that uses a deterministic ordering guarantee to significantly reduce the normally prohibitive contention costs associated with distributed transactions. Unlike previous deterministic database system prototypes, Calvin supports disk-based storage, scales near-linearly on a cluster of commodity machines, and has no single point of failure. By replicating transaction inputs rather than effects, Calvin is also able to support multiple consistency levels---including Paxos-based strong consistency across geographically distant replicas---at no cost to transactional throughput.

Advanced Partitioning Techniques for Massively Distributed Computation

Jingren Zhou, Microsoft; Nicolás Bruno, Microsoft; Wei Lin, Microsoft

An increasing number of companies rely on distributed data storage and processing over large clusters of commodity machines for critical business decisions. Although plain MapReduce systems provide several benefits, they carry certain limitations that impact developer productivity and optimization opportunities. Higher level programming languages plus conceptual data models have recently emerged to address such limitations. These languages offer a single machine programming abstraction and are able to perform sophisticated query optimization and apply efficient execution strategies.
In massively distributed computation, data shuffling is typically the most expensive operation and can lead to serious performance bottlenecks if not done properly. An important optimization opportunity in this environment is that of judicious placement of repartitioning operators and choice of alternative implementations. In this paper we discuss advanced partitioning strategies, their implementation, and how they are integrated in the Microsoft Scope system. We show experimentally that our approach significantly improves performance for a large class of real-world jobs.

SkewTune: Mitigating Skew in MapReduce Applications

YongChul Kwon, University of Washington; Magdalena Balazinska, University of Washington; Bill Howe, University of Washington; Jerome Rolia, HP Labs

We present an automatic skew mitigation approach for user-defined MapReduce programs and present SkewTune, a system that implements this approach as a drop-in replacement for an existing MapReduce implementation. There are three key challenges: (a) require no extra input from the user yet work for all MapReduce applications, (b) be completely transparent, and (c) impose minimal overhead if there is no skew. The SkewTune approach addresses these challenges and works as follows: When a node in the cluster becomes idle, SkewTune identifies the task with the greatest expected remaining processing time. The unprocessed input data of this straggling task is then proactively repartitioned in a way that fully utilizes the nodes in the cluster and preserves the ordering of the input data so that the original output can be reconstructed by concatenation. We implement SkewTune as an extension to Hadoop and evaluate its effectiveness using several real applications. The results show that SkewTune can significantly reduce job runtime in the presence of skew and adds little to no overhead in the absence of skew.

SIGMOD Research 2: Indexing and Physical Database Design I

Parallel Main-Memory Indexing for Moving-Object Query and Update Workloads

Darius Sidlauskas, Aalborg University; Simonas Saltenis, Aalborg University; Christian S. Jensen, Aarhus University

We are witnessing a proliferation of Internet-worked, geo-positioned mobile devices such as smartphones and personal navigation devices. Likewise, location-related services that target the users of such devices are proliferating. Consequently, server-side infrastructures are needed that are capable of supporting the location-related query and update workloads generated by very large populations of such moving objects.
This paper presents a main-memory indexing technique that aims to support such workloads. The technique, called PGrid, uses a grid structure that is capable of exploiting the parallelism offered by modern processors. Unlike earlier proposals that maintain separate structures for updates and queries, PGrid allows both long-running queries and rapid updates to operate on a single data structure and thus offers up-to-date query results. Because PGrid does not rely on creating snapshots, it avoids the stop-the-world problem that occurs when workload processing is interrupted to perform such snapshotting. Its concurrency control mechanism relies instead on hardware-assisted atomic updates as well as object-level copying, and it

treats updates as non-divisible operations rather than as combinations of deletions and insertions; thus, the query semantics guarantee that no objects are missed in query results.
Empirical studies demonstrate that PGrid scales near-linearly with the number of hardware threads on four modern multi-core processors. Since both updates and queries are processed on the same current data-store state, PGrid outperforms snapshot-based techniques in terms of both query freshness and CPU cycle-wise efficiency.

Divergent Physical Design Tuning for Replicated Databases

Mariano P. Consens, University of Toronto; Kleoni Ioannidou, University of California, Santa Cruz; Jeff LeFevre, University of California, Santa Cruz; Neoklis Polyzotis, University of California, Santa Cruz

We introduce divergent designs as a novel tuning paradigm for database systems that employ replication. A divergent design installs a different physical configuration (e.g., indexes and materialized views) with each database replica, specializing replicas for different subsets of the workload. At runtime, queries are routed to the subset of the replicas configured to yield the most efficient execution plans. When compared to uniformly designed replicas, divergent replicas can potentially execute their subset of the queries significantly faster, and their physical configurations could be initialized and maintained (updated) in less time. However, the specialization of divergent replicas limits the ability to load-balance the workload at runtime.
We formalize the divergent design problem, characterize the properties of good designs, and analyze the complexity of identifying the optimal divergent design. Our paradigm captures the trade-off between load balancing among all n replicas vs. load balancing among m <= n specialized replicas. We develop an effective algorithm (leveraging single-node-tuning functionality) to compute good divergent designs for all the points of this trade-off. Experimental results validate the effectiveness of the algorithm and demonstrate that divergent designs can substantially improve workload performance.

Skew-Aware Automatic Database Partitioning in Shared-Nothing, Parallel OLTP Systems

Andrew Pavlo, Brown University; Carlo Curino, Yahoo! Research; Stanley Zdonik, Brown University

The advent of affordable, shared-nothing computing systems portends a new class of parallel database management systems (DBMS) for on-line transaction processing (OLTP) applications that scale without sacrificing ACID guarantees. The performance of these DBMSs is predicated on the existence of an optimal database design that is tailored for the unique characteristics of OLTP workloads. Deriving such designs for modern DBMSs is difficult, especially for enterprise-class} OLTP systems, since they impose extra challenges: the use of stored procedures, the need for load balancing in the presence of time-varying skew, complex schemas, and deployments with larger number of partitions.
To this purpose, we present a novel approach to automatically partitioning databases for enterprise-class OLTP systems that significantly extends the state of the art by: (1) minimizing the number distributed transactions, while concurrently mitigating the effects of temporal skew in both the data distribution and accesses, (2) extending the design space to include replicated secondary indexes, (4) organically handling stored procedure routing, and (3) scaling of schema complexity, data size, and number of partitions. This effort builds on two key technical contributions: an analytical cost model that can be used to quickly estimate the relative coordination cost and skew for a given workload and a candidate database design, and an informed exploration of the huge solution space based on large neighborhood search. To evaluate our methods, we integrated our database design tool with a high-performance parallel, main memory DBMS and compared our methods against both popular heuristics and a state-of-the-art research prototype. Using a diverse set of benchmarks, we show that our approach improves throughput by up to a factor of 16x over these other approaches.

SIGMOD Research 3: Data Cleaning and Integration

Sample-Driven Schema Mapping

Li Qian, University of Michigan, Ann Arbor; Michael J. Cafarella, University of Michigan, Ann Arbor; H. V. Jagadish, University of Michigan, Ann Arbor

End-users increasingly find the need to perform light-weight, customized schema mapping. State-of-the-art tools provide powerful functions to generate schema mappings, but they usually require an in-depth understanding of the semantics of multiple schemas and their correspondences, and are thus not suitable for users who are technically unsophisticated or when a large number of mappings must be performed.
We propose a system for *sample-driven* schema mapping. It automatically constructs schema mappings, in real time, from user-input sample target instances. Because the user does not have to provide any explicit attribute-level match information, she is isolated from the possibly complex structure and semantics of both the source schemas and the mappings. In addition, the user never has to master any operations specific to schema mappings: she simply types data values into a spreadsheet-style interface. As a result, the user can construct mappings with a much lower cognitive burden.
In this paper, we present Mweaver, a prototype sample-driven schema mapping system. It employs novel algorithms that enable the system to obtain desired mapping results while meeting interactive response performance requirements. We show the results of a user study that compares Mweaver with two state-of-the-art mapping tools across several mapping

tasks, both real and synthetic. These suggest that the Mweaver system enables users to perform practical mapping tasks in about *1/5th* the time needed by the state-of-the-art tools.

Can We Beat the Prefix Filtering? An Adaptive Framework for Similarity Join and Search

Jiannan Wang, Tsinghua University; Guoliang Li, Tsinghua University; Jianhua Feng, Tsinghua University

As two important operations in data cleaning, similarity join and similarity search have attracted much attention recently. Existing methods to support similarity join usually adopt a prefix-filtering-based framework. They select a prefix of each object and prune object pairs whose prefixes have no overlap. We have an observation that prefix lengths have significant effect on the performance. Different prefix lengths lead to significantly different performance, and prefix filtering does not always achieve high performance. To address this problem, in this paper we propose an adaptive framework to support similarity join. We propose a cost model to judiciously select an appropriate prefix for each object. To efficiently select prefixes, we devise effective indexes. We extend our method to support similarity search. Experimental results show that our framework beats the prefix-filtering-based framework and achieves high efficiency.

InfoGather: Entity Augmentation and Attribute Discovery by Holistic Matching with Web Tables

Mohamed Yakout, Purdue University; Kris Ganjam, Microsoft Research; Kaushik Chakrabarti, Microsoft Research; Surajit Chaudhuri, Microsoft Research

The Web contains a vast corpus of HTML tables, specifically entity-attribute tables. We present three core operations, namely entity augmentation by attribute name, entity augmentation by example and attribute discovery, that are useful for ``information gathering'' tasks (e.g., researching for products or stocks). We propose to use web table corpus to perform them automatically. We require the operations to have high precision and coverage, have fast (ideally interactive) response times and be applicable to any arbitrary domain of entities. The naive approach that attempts to directly match the user input with the web tables suffers from poor precision and coverage.

Our key insight is that we can achieve much higher precision and coverage by considering indirectly matching tables in addition to the directly matching ones. The challenge is to be robust to spuriously matched tables: we address it by developing a holistic matching framework based on topic sensitive pagerank and an augmentation framework that aggregates predictions from multiple matched tables. We propose a novel architecture that leverages preprocessing in MapReduce to achieve extremely fast response times at query time. Our experiments on real-life datasets and 573M web tables show that our approach has (i) significantly higher precision and coverage and (ii) four orders of magnitude faster response times compared with the state-of-the-art approach.

SIGMOD Research 4: Query Processing and Optimization

Interactive Regret Minimization

Danupon Nanongkai, University of Vienna; Ashwin Lall, Denison University; Atish Das Sarma, Google Research; Kazuhisa Makino, University of Tokyo

We study the notion of *regret ratio* proposed in Nanongkai et al. to deal with multi-criteria decision making in database systems. The regret minimization query proposed in Nanongkai et al. was shown to have features of both skyline and top-*k*: it does not need information from the user but still controls the output size. While this approach is suitable for obtaining a reasonably small regret ratio, it is still open whether one can make the regret ratio arbitrarily small. Moreover, it remains open whether reasonable questions can be asked to the users in order to improve efficiency of the process.

In this paper, we study the problem of minimizing regret ratio when the system is enhanced with *interaction*. We assume that when presented with a set of tuples the user can tell which tuple is most preferred. Under this assumption, we develop the problem of *interactive regret minimization* where we fix the number of questions and tuples per question that we can display, and aim at minimizing the regret ratio. We try to answer two questions in this paper: (1) How much does interaction help? That is, how much can we improve the regret ratio when there are interactions? (2) How efficient can interaction be? In particular, we measure how many questions we have to ask the user in order to make her regret ratio small enough.

We answer both questions from both theoretical and practical standpoints. For the first question, we show that interaction can reduce the regret ratio almost *exponentially*. To do this, we prove a lower bound for the previous approach (thereby resolving an open problem from Nanongkai et al.), and develop an almost-optimal upper bound that makes the regret ratio exponentially smaller. Our experiments also confirm that, in practice, interactions help in improving the regret ratio by many orders of magnitude. For the second question, we prove that when our algorithm shows a reasonable number of points per question, it only needs a few questions to make the regret ratio small. Thus, interactive regret minimization seems to be a necessary and sufficient way to deal with multi-criteria decision making in database systems.

MCJoin: A Memory-Constrained Join for Column-Store Main-Memory Databases

Steven Keith Begley, La Trobe University; Zhen He, La Trobe University; Yi-Ping Phoebe Chen, La Trobe University

There exists a need for high performance, read-only main-memory database systems for OLAP-style application scenarios. Most of the existing works in this area are centered around the domain of column-store databases, which are particularly well suited to OLAP-style scenarios and have been shown to overcome the memory bottleneck issues that have been found to hinder the more traditional row-store database systems. One of the main database operations these systems are focused on optimizing is the JOIN operation. However, all these existing systems use join algorithms that are designed with the unrealistic assumption that there is unlimited temporary memory available to perform the join. In contrast, we propose a Memory Constrained Join algorithm (MCJoin) which is both high performing and also performs all of its operations within a tight given memory constraint. Extensive experimental results show that MCJoin outperforms a naive memory constrained version of the state-of-the-art Radix-Clustered Hash Join algorithm in all of the situations tested, with margins of up to almost 500%.

Holistic Optimization by Prefetching Query Results

Karthik Ramachandra, Indian Institute of Technology Bombay; S. Sudarshan, Indian Institute of Technology Bombay

In this paper we address the problem of optimizing performance of database/web-service backed applications by means of automatically prefetching query results. Prefetching has been performed in earlier work based on predicting query access patterns; however such prediction is often of limited value, and can perform unnecessary prefetches. There has been some earlier work on program analysis and rewriting to automatically insert prefetch requests; however, such work has been restricted to rewriting of single procedures. In many cases, the query is in a procedure which does not offer much scope for prefetching within the procedure; in contrast, our approach can perform prefetching in a calling procedure, even when the actual query is in a called procedure, thereby greatly improving the benefits due to prefetching. Our approach does not perform any intrusive changes to the source code, and places prefetch instructions at the earliest possible points while avoiding wasteful prefetches. We have incorporated our techniques into a tool for holistic optimization called DBridge, to prefetch query results in Java programs that use JDBC. Our tool can be easily extended to handle Hibernate API calls as well as Web service requests. Our experiments on several real world applications demonstrate the applicability and significant performance gains due to our techniques.

SIGMOD Research 5: Social Networks and Graph Databases I

Managing Large Dynamic Graphs Efficiently

Jayanta Mondal, University of Maryland; Amol Deshpande, University of Maryland

There is an increasing need to ingest, manage, and query large volumes of graph-structured data arising in applications like social networks, communication networks, biological networks, and so on. Graph databases that can explicitly reason about the graphical nature of the data, that can support flexible schemas and node-centric or edge-centric analysis and querying, are ideal for storing such data. However, although there is much work on single-site graph databases and on efficiently executing different types of queries over large graphs, to date there is little work on understanding the challenges in distributed graph databases, needed to handle the large scale of such data. In this paper, we propose the design of an in-memory, distributed graph data management system aimed at managing a large-scale dynamically changing graph, and supporting low-latency query processing over it. The key challenge in a distributed graph database is that, partitioning a graph across a set of machines inherently results in a large number of distributed traversals across partitions to answer even simple queries. We propose aggressive replication of the nodes in the graph for supporting low-latency querying, and investigate three novel techniques to minimize the communication bandwidth and the storage requirements. First, we develop a hybrid replication policy that monitors node read-write frequencies to dynamically decide what data to replicate, and whether to do "eager" or "lazy" replication. Second, we propose a clustering-based approach to amortize the costs of making these replication decisions. Finally, we propose using a "fairness" criterion to dictate how replication decisions should be made. We provide both theoretical analysis and efficient algorithms for the optimization problems that arise. We have implemented our framework as a middleware on top of the open-source CouchDB key-value store. We evaluate our system on a social graph, and show that our system is able to handle very large graphs efficiently, and that it reduces the network bandwidth consumption significantly.

Query Preserving Graph Compression

Wenfei Fan, University of Edinburgh; Jianzhong Li, Harbin Institute of Technology; Xin Wang, University of Edinburgh; Yinghui Wu, University of Edinburgh and University of California, Santa Barbara

It is common to find graphs with millions of nodes and billions of edges in, e.g., social networks. Queries on such graphs are often prohibitively expensive. These motivate us to propose query preserving graph compression, to compress graphs relative to a class {\cal Q} of queries of users' choice. We compute a small G_r from a graph G such that (a) for any query Q \in {\cal Q}, Q(G) = Q'(G_r), where Q' \in {\cal Q} can be efficiently computed from Q; and (b) any algorithm for computing Q(G) can be directly applied to evaluating Q' on G_r as is. That is, while we cannot lower the complexity of evaluating graph queries, we reduce data graphs while preserving the answers to all the queries in {\cal Q}. To verify the effectiveness of this approach, (1) we develop compression strategies for two classes of queries: reachability and graph pattern queries via (bounded) simulation. We show that graphs can be efficiently compressed via a reachability equivalence relation and graph bisimulation, respectively, while reserving query answers. (2) We provide techniques for aintaining compressed graph G_r in response to changes {\Delta G} to the original graph G. We show that the incremental maintenance problems are unbounded for the two lasses of queries, i.e., their costs are not a function of the size of {\Delta G} and changes in G_r. Nevertheless, we develop incremental algorithms that depend only on {\Delta G} and G_r, independent of G, i.e., we do not have to decompress G_r to propagate the changes. (3) Using real-life data, we experimentally verify that our compression techniques could reduce graphs in average by 95% for reachability and 57% for graph pattern matching, and that our incremental maintenance algorithms are efficient.

SCARAB: Scaling Reachability Computation on Large Graphs

Ruoming Jin, Kent State University; Ning Ruan, Kent State University; Saikat Dey, Kent State University; Jeffrey Xu Yu, The Chinese University of Hong Kong

Most of the existing reachability indices perform well on small- to medium- size graphs, but reach a scalability bottleneck around one million vertices/edges. As graphs become increasingly large, scalability is quickly becoming the major research challenge for the reachability computation today. Can we construct indices which scale to graphs with tens of millions of vertices and edges? Can the existing reachability indices which perform well on moderate-size graphs be scaled to very large graphs? In this paper, we propose *SCARAB* (standing for SCAlable ReachABility), a unified reachability computation framework: it not only can scale the existing state-of-the-art reachability indices, which otherwise could only be constructed and work on moderate size graphs, but also can help speed up the online query answering approaches. Our experimental results demonstrate that SCARAB can perform on graphs with millions of vertices/edges and is also much faster then GRAIL, the state-of-the-art scalability index approach.

SIGMOD Research 6: Data Visualization, Error Reporting

Skimmer: Rapid Scrolling of Relational Query Results

Manish Singh, University of Michigan, Ann Arbor; Arnab Nandi, Ohio State University; H. V. Jagadish, University of Michigan, Ann Arbor

A relational database often yields a large set of tuples as the result of a query. Users browse this result set to find the information they require. If the result set is large, there may be many pages of data to browse. Since results comprise tuples of alphanumeric values that have few visual markers, it is hard to browse the data quickly, even if it is sorted.
In this paper, we describe the design of a system for browsing relational data by scrolling through it at a high speed. Rather than showing the user a fast-changing blur, the system presents the user with a small number of representative tuples. Representative tuples are selected to provide a "good impression" of the query result. We show that the information loss to the user is limited, even at high scrolling speeds, and that our algorithms can pick good representatives fast enough to provide for real-time, high-speed scrolling over large datasets.

Efficient Spatial Sampling of Large Geographical Tables

Anish Das Sarma, Google; Hongrae Lee, Google; Hector Gonzalez, Google; Jayant Madhavan, Google; Alon Halevy, Google

Large-scale map visualization systems play an increasingly important role in presenting geographic datasets to end users. Since these datasets can be extremely large, a map rendering system often needs to select a small fraction of the data to visualize them in a limited space. This paper addresses the fundamental challenge of thinning: determining appropriate samples of data to be shown on specific geographical regions and zoom levels. Other than the sheer scale of the data, the thinning problem is challenging because of a number of other reasons: (1) data can consist of complex geographical shapes, (2) rendering of data needs to satisfy certain constraints, such as data being preserved across zoom levels and adjacent regions, and (3) after satisfying the constraints, an optimal solution needs to be chosen based on objectives such as maximality, fairness, and importance of data.
This paper formally defines and presents a complete solution to the thinning problem. First, we express the problem as a

integer programming formulation that efficiently solves thinning for desired objectives. Second, we present more efficient solutions for maximality, based on DFS traversal of a spatial tree. Third, we consider the common special case of point datasets, and present an even more efficient randomized algorithm. Finally, we have implemented all techniques from this paper in Google Maps visualizations of Fusion Tables, and we describe a set of experiments that demonstrate the tradeoffs among the algorithms.

Declarative Error Management for Robust Data-Intensive Applications

Carl-Christian Kanne, Platfora Inc.; Vuk Ercegovac, IBM Almaden Research Center

We present an approach to declaratively manage run-time errors in data-intensive applications. When large volumes of raw data meet complex third-party libraries, deterministic run-time errors become likely, and existing query processors typically stop without returning a result when a run-time error occurs. The ability to degrade gracefully in the presence of run-time errors, and partially execute jobs, is typically limited to specific operators such as bulkloading.
We generalize this concept to all operators of a query processing system, introducing a novel data type ``partial result with errors" and corresponding operators. We show how to extend existing error-unaware operators to support this type, and as an added benefit, eliminate side-effect based error reporting. We use declarative specifications of acceptable results to control the semantics of error-aware operators.
We have incorporated our approach into a declarative query processing system, which compiles the language constructs into instrumented execution plans for clusters of machines. We experimentally validate that the instrumentation overhead is below 20% in microbenchmarks, and not detectable when running I/O-intensive workloads.

SIGMOD Research 7: Storage Systems, Query Processing and Optimization

bLSM: A General Purpose Log Structured Merge Tree

Russell Sears, Yahoo! Research; Raghu Ramakrishnan, Yahoo! Research

Data management workloads are increasingly write-intensive and subject to strict latency SLAs. This presents a dilemma: Update in place systems have unmatched latency but poor write throughput. In contrast, existing log structured techniques improve write throughput but sacrifice read performance and exhibit unacceptable latency spikes.We begin by presenting a new performance metric: read fanout, and argue that, with read and write amplification, it better characterizes real-world indexes than approaches such as asymptotic analysis and price/performance.
We then present bLSM, a Log Structured Merge (LSM) tree with the advantages of B-Trees and log structured approaches: (1) Unlike existing log structured trees, bLSM has near-optimal read and scan performance, and (2) its new ``spring and gear" merge scheduler bounds write latency without impacting throughput or allowing merges to block writes for extended periods of time. It does this by ensuring merges at each level of the tree make steady progress without resorting to techniques that degrade read performance.
We use Bloom filters to improve index performance, and find a number of subtleties arise. First, we ensure reads can stop after finding one version of a record. Otherwise, frequently written items would incur multiple B-Tree lookups. Second, many applications check for existing values at insert. Avoiding the seek performed by the check is crucial.

Skeleton Automata for FPGAs: Reconfiguring without Reconstructing

Jens Teubner, ETH Zürich; Louis Woods, ETH Zürich; Chongling Nie, ETH Zürich

While the performance opportunities of field-programmable gate arrays (FPGAs) for high-volume query processing are well-known, system makers still have to compromise between desired query expressiveness and high compilation effort. The cost of the latter is the primary limitation in building efficient FPGA/CPU hybrids.
In this work we report on an FPGA-based stream processing engine that does not have this limitation. We provide a hardware implementation of XML projection that can be reconfigured in less than a micro-second, yet supports a rich and expressive dialect of XPath. By performing XML projection in the network, we can fully leverage its filtering effect and improve XQuery performance by several factors.
These improvements are made possible by a new design approach for FPGA acceleration, called skeleton automata. Skeleton automata separate the structure of finite-state automata from their semantics. Since individual queries only affect the latter, with our approach query workload changes can be accommodated fast and with high expressiveness.

NoDB: Efficient Query Execution on Raw Data Files

Ioannis Alagiannis, EPFL; Renata Borovica, EPFL; Miguel Branco, EPFL; Stratos Idreos, CWI; Anastasia Ailamaki, EPFL

As data collections become larger and larger, data loading evolves to a major bottleneck. Many applications already avoid using database systems, e.g., scientific data analysis and social networks, due to the complexity and the increased data-to-query time. For such applications data collections keep growing fast, even on a daily basis, and we are already in the era of data deluge where we have much more data than what we can move, store, let alone analyze.

Our contribution in this paper is the design and roadmap of a new paradigm in database systems, called NoDB, which do not require data loading while still maintaining the whole feature set of a modern database system. In particular, we show how to make raw data files a first-class citizen, fully integrated with the query engine. Through our design and lessons learned by implementing the NoDB philosophy over a modern DBMS, we discuss the fundamental limitations as well as the strong opportunities that such a research path brings. We identify performance bottlenecks specific for in situ processing, namely the repeated parsing and tokenizing overhead and the expensive data type conversion costs. To address these problems, we introduce an adaptive indexing mechanism that maintains positional information to provide efficient access to raw data files, together with a flexible caching structure.
Our implementation over PostgreSQL, called PostgresRaw, is able to avoid the loading cost completely, while matching the query performance of plain PostgreSQL and even outperforming it in many cases. We conclude that NoDB systems are feasible to design and implement over modern database architectures, bringing an unprecedented positive effect in usability and performance.

SIGMOD Research 8: Data Streams and Sensor Networks

High-Performance Complex Event Processing over XML Streams

Barzan Mozafari, University of California, Los Angeles; Kai Zeng, University of California, Los Angeles; Carlo Zaniolo, University of California, Los Angeles

Much research attention has been given to delivering high-performance systems that are capable of complex event processing (CEP) in a wide range of applications. However, many current CEP systems focus on processing efficiently data having a simple structure, and are otherwise limited in their ability to support efficiently complex continuous queries on structured or semi-structured information. However, XML streams represent a very popular form of data exchange, comprising large portions of social network and RSS feeds, financial records, configuration files, and similar applications requiring advanced CEP queries. In this paper, we present the XSeq language and system that support CEP on XML streams, via an extension of XPath that is both powerful and amenable to an efficient implementation. Specifically, the XSeq language extends XPath with natural operators to express sequential and Kleene-* patterns over XML streams, while remaining highly amenable to efficient implementation. XSeq is designed to take full advantage of recent advances in the field of automata on Visibly Pushdown Automata (VPA), where higher expressive power can be achieved without compromising efficiency (whereas the amenability to efficient implementation was not demonstrated in XPath extensions previously proposed).
We illustrate XSeq's power for CEP applications through examples from different domains, and provide formal results on its expressiveness and complexity. Finally, we present several optimization techniques for XSeq queries. Our extensive experiments indicate that XSeq brings outstanding performance to CEP applications: two orders of magnitude improvement are obtained over the same queries executed in general-purpose XML engines.

Prediction-Based Geometric Monitoring over Distributed Data Streams

Nikos Giatrakos, University of Piraeus; Antonios Deligiannakis, Technical University of Crete; Minos Garofalakis, Technical University of Crete; Izchak Sharfman, Technion; Assaf Schuster, Technion

Many modern streaming applications, such as online analysis of financial, network, sensor and other forms of data are inherently distributed in nature. An important query type that is the focal point in such application scenarios regards actuation queries, where proper action is dictated based on a trigger condition placed upon the current value that a monitored function receives. Recent work studies the problem of (non-linear) sophisticated function tracking in a distributed manner. The main concept behind the geometric monitoring approach proposed there, is for each distributed site to perform the function monitoring over an appropriate subset of the input domain. In the current work, we examine whether the distributed monitoring mechanism can become more efficient, in terms of the number of communicated messages, by extending the geometric monitoring framework to utilize prediction models. We initially describe a number of local estimators (predictors) that are useful for the applications that we consider and which have already been shown particularly useful in past work. We then demonstrate the feasibility of incorporating predictors in the geometric monitoring framework and show that prediction-based geometric monitoring in fact generalizes the original geometric monitoring framework. We propose a large variety of different prediction-based monitoring models for the distributed threshold monitoring of complex functions. Our extensive experimentation with a variety of real data sets, functions and parameter settings indicates that our approaches can provide significant communication savings ranging between two times and up to three orders of magnitude, compared to the transmission cost of the original monitoring framework.

Online Windowed Subsequence Matching over Probabilistic Sequences

Zheng Li, University of Massachusetts, Lowell; Tingjian Ge, University of Massachusetts, Lowell

Windowed subsequence matching over deterministic strings has been studied in previous work in the contexts of knowledge discovery, data mining, and molecular biology. However, we observe that in these applications, as well as in data stream monitoring, complex event processing, and time series data processing in which streams can be mapped to strings, the

strings are often noisy and probabilistic. We study this problem in the online setting where efficiency is paramount. We first formulate the query semantics, and propose an exact algorithm. Then we propose a randomized approximation algorithm that is faster and, in the mean time, provably accurate. Moreover, we devise a filtering algorithm to further enhance the efficiency with an optimization technique that is adaptive to sequence stream contents. Finally, we propose algorithms for patterns with negations. In order to verify the algorithms, we conduct a systematic empirical study using three real datasets and some synthetic datasets.

SIGMOD Research 9: Mobile Databases

MaskIt: Privately Releasing User Context Streams for Personalized Mobile Applications

Michaela Goetz, Twitter; Suman Nath, Microsoft Research; Johannes Gehrke, Cornell University

The rise of smartphones equipped with various sensors has enabled personalization of various applications based on user contexts extracted from sensor readings. At the same time it has raised serious concerns about the privacy of user contexts.
In this paper, we present MASKIT, a technique to filter a user context stream that provably preserves privacy. The filtered context stream can be released to applications or be used to answer their queries. Privacy is defined with respect to a set of sensitive contexts specified by the user. MASKIT limits what adversaries can learn from the filtered stream about the user being in a sensitive context - even if the adversaries are powerful and have knowledge about the filtering system and temporal correlations in the context stream.
At the heart of MASKIT is a privacy check deciding whether to release or suppress the current user context. We present two novel privacy checks and explain how to choose the one with the higher utility for a user. Our experiments on real smartphone context traces of 91 users demonstrate the high utility of MASKIT.

Authenticating Location-Based Services without Compromising Location Privacy

Haibo Hu, Hong Kong Baptist University; Jianliang Xu, Hong Kong Baptist University; Qian Chen, Hong Kong Baptist University; Ziwei Yang, Hong Kong Baptist University

The popularity of mobile social networking services (mSNSs) is propelling more and more businesses, especially those in retailing and marketing, into mobile and location-based forms. To address the trust issue, the service providers are expected to deliver their location-based services in an authenticatable manner, so that the correctness of the service results can be verified by the client. However, existing works on query authentication cannot preserve the privacy of the data being queried, which are sensitive user locations when it comes to location-based services and mSNSs. In this paper, we address this challenging problem by proposing a comprehensive solution that preserves unconditional location privacy when authenticating range queries. Three authentication schemes for *R*-tree and grid-file index, together with two optimization techniques, are developed. Cost models, security analysis, and experimental results consistently show the effectiveness, reliability and robustness of the proposed schemes under various system settings and query workloads.

Effective Caching of Shortest Paths for Location-Based Services

Jeppe Rishede Thomsen, Hong Kong Polytechnic University; Man Lung Yiu, Hong Kong Polytechnic University; Christian S. Jensen, Aarhus University

Web search is ubiquitous in our daily lives. Caching has been extensively used to reduce the computation time of the search engine and reduce the network traffic beyond a proxy server. Another form of web search, known as online shortest path search, is popular due to advances in geo-positioning. However, existing caching techniques are ineffective for shortest path queries.This is due to several crucial differences between web search results and shortest path results, in relation to query matching, cache item overlapping, and query cost variation.
Motivated by this, we identify several properties that are essential to the success of effective caching for shortest path search. Our cache exploits the optimal subpath property, which allows a cached shortest path to answer any query with source and target nodes on the path. We utilize statistics from query logs to estimate the benefit of caching a specific shortest path, and we employ a greedy algorithm for placing beneficial paths in the cache. Also, we design a compact cache structure that supports efficient query matching at runtime. Empirical results on real datasets confirm the effectiveness of our proposed techniques.

SIGMOD Research 10: Data Analytics

Towards a Unified Architecture for in-RDBMS Analytics

Xixuan Feng, University of Wisconsin-Madison; Arun Kumar, University of Wisconsin-Madison; Benjamin Recht, University of Wisconsin-Madison; Christopher Ré, University of Wisconsin-Madison

The increasing use of statistical data analysis in enterprise applications has created an arms race among database vendors to offer ever more sophisticated in-database analytics. One challenge in this race is that each new statistical technique must be implemented from scratch in the RDBMS, which leads to a lengthy and complex development process. We argue that the

root cause for this overhead is the lack of a unified architecture for in-database analytics. Our main contribution in this work is to take a step towards such a unified architecture. A key benefit of our unified architecture is that performance optimizations for analytics techniques can be studied generically instead of an ad hoc, per-technique fashion. In particular, our technical contributions are theoretical and empirical studies of two key factors that we found impact performance: the order data is stored, and parallelization of computations on a single-node multicore RDBMS. We demonstrate the feasibility of our architecture by integrating several popular analytics techniques into two commercial and one open-source RDBMS. Our architecture requires changes to only a few dozen lines of code to integrate a new statistical technique. We then compare our approach with the native analytics tools offered by the commercial RDBMSes on various analytics tasks, and validate that our approach achieves competitive or higher performance, while still achieving the same quality.

Tiresias: The Database Oracle for How-To Queries

Alexandra Meliou, University of Washington; Dan Suciu, University of Washington

How-To queries answer fundamental data analysis questions of the form: ``How should the input change in order to achieve the desired output". As a \Reverse Data Management problem, the evaluation of how-to queries is harder than their ``forward" counterpart: hypothetical, or what-if queries.

In this paper, we present Tiresias, the first system that provides support for how-to queries, allowing the definition and integrated evaluation of a large set of constrained optimization problems, specifically Mixed Integer Programming problems, on top of a relational database system. Tiresias generates the problem variables, constraints and objectives by issuing standard SQL statements, allowing for its integration with any RDBMS.

The contributions of this work are the following: (a) we define how-to queries using possible world semantics, and propose the specification language *TiQL* (for Tiresias Query Language) based on simple extensions to standard Datalog. (b) We define translation rules that generate a Mixed Integer Program (MIP) from TiQL specifications, which can be solved using existing tools. (c) Tiresias implements powerful ``data-aware" optimizations that are beyond the capabilities of modern MIP solvers, dramatically improving the system performance. (d) Finally, an extensive performance evaluation on the TPC-H dataset demonstrates the effectiveness of these optimizations, particularly highlighting the ability to apply divide-and-conquer methods to break MIP problems into smaller instances.

GUPT: Privacy Preserving Data Analysis Made Easy

Prashanth Mohan, University of California, Berkeley; Abhradeep Thakurta, Pennsylvania State University; Elaine Shi, University of California, Berkeley; Dawn Song, University of California, Berkeley; David Culler, University of California, Berkeley

It is often highly valuable for organizations to have their data analyzed by external agents. However, any program that computes on potentially sensitive data risks leaking information through its output. Differential privacy provides a theoretical framework for processing data while protecting the privacy of individual records in a dataset. Unfortunately, it has seen limited adoption because of the loss in output accuracy, the difficulty in making programs differentially private, lack of mechanisms to describe the privacy budget in a programmer's utilitarian terms, and the challenging requirement that data owners and data analysts manually distribute the limited privacy budget between queries.

This paper presents the design and evaluation of a new system, GUPT, that overcomes these challenges. Unlike existing differentially private systems such as PINQ and Airavat, it guarantees differential privacy to programs not developed with privacy in mind, makes no trust assumptions about the analysis program, and is secure to all known classes of side-channel attacks.

GUPT uses a new model of data sensitivity that degrades privacy of data over time. This enables efficient allocation of different levels of privacy for different user applications while guaranteeing an overall constant level of privacy and maximizing the utility of each application. GUPT also introduces techniques that improve the accuracy of output while achieving the same level of privacy. These approaches enable GUPT to easily execute a wide variety of data analysis programs while providing both utility and privacy.

SIGMOD Research 11: Crowdsourcing, Uncertainty in Databases

CrowdScreen: Algorithms for Filtering Data with Humans

Aditya G. Parameswaran, Stanford University; Héctor García-Molina, Stanford University; Hyunjung Park, Stanford University; Neoklis Polyzotis, University of California, Santa Cruz; Aditya Ramesh, Stanford University; Jennifer Widom, Stanford University

Given a large set of data items, we consider the problem of filtering them based on a set of properties that can be verified by humans. This problem is commonplace in crowdsourcing applications, and yet, to our knowledge, no one has considered the formal optimization of this problem. (Typical solutions use heuristics to solve the problem.) We formally state a few different variants of this problem. We develop deterministic and probabilistic algorithms to optimize the expected cost (i.e., number of questions) and expected error. We experimentally show that our algorithms provide definite gains with respect to

other strategies. Our algorithms can be applied in a variety of crowdsourcing scenarios and can form an integral part of any query processor that uses human computation.

Local Structure and Determinism in Probabilistic Databases

Theodoros Rekatsinas, University of Maryland; Amol Deshpande, University of Maryland; Lise Getoor, University of Maryland

While extensive work has been done on evaluating queries over tuple-independent probabilistic databases, query evaluation over correlated data has received much less attention even though the support for correlations is essential for many natural applications of probabilistic databases, e.g., information extraction, data integration, computer vision, etc. In this paper, we develop a novel approach for efficiently evaluating probabilistic queries over correlated databases where correlations are represented using a "factor graph", a class of graphical models widely used for capturing correlations and performing statistical inference. Our approach exploits the specific values of the factor parameters and the determinism in the correlations, collectively called "local structure", to reduce the complexity of query evaluation. Our framework is based on "arithmetic circuits", factorized representations of probability distributions that can exploit such local structure. Traditionally, arithmetic circuits are generated following a compilation process and can not be updated directly. We introduce a generalization of arithmetic circuits, called "annotated arithmetic circuits", and a novel algorithm for updating them, which enables us to answer probabilistic queries efficiently. We present a comprehensive experimental analysis and show speed-ups of at least one order of magnitude in many cases.

So Who Won? Dynamic Max Discovery with the Crowd

Stephen Guo, Stanford University; Aditya G. Parameswaran, Stanford University; Héctor García-Molina, Stanford University

We consider a crowdsourcing database system that may cleanse, populate, or filter its data by using human workers. Just like a conventional DB system, such a crowdsourcing DB system requires data manipulation functions such as select, aggregate, maximum, average, and so on, except that now it must rely on human operators (that for example compare two objects) with very different latency, cost and accuracy characteristics. In this paper, we focus on one such function, maximum, that finds the highest ranked object or tuple in a set. In particularm we study two problems: given a set of votes (pairwise comparisons among objects), how do we select the maximum? And how do we improve our estimate by requesting additional votes? We show that in a crowdsourcing DB system, the optimal solution to both problems is NP-Hard. We then provide heuristic functions to select the maximum given evidence, and to select additional votes. We experimentally evaluate our functions to highlight their strengths and weaknesses.

SIGMOD Research 12: Top-k Query Processing and Optimization

Processing a Large Number of Continuous Preference Top-k Queries

Albert Yu, Duke University; Pankaj K. Agarwal, Duke University; Jun Yang, Duke University

Given a set of objects, each with multiple numeric attributes, a (preference) top-k query retrieves the k objects with the highest scores according to a user preference, defined as a linear combination of attribute values. We consider the problem of processing a large number of continuous top-k queries, each with its own preference. When objects or user preferences change, the query results must be updated. We present a dynamic index that supports the reverse top-k query, which is of independent interest. Combining this index with another one for top-k queries, we develop a scalable solution for processing many continuous top-k queries that exploits the clusteredness in user preferences. We also define an approximate version of the problem and present a solution significantly more efficient than the exact one with little loss in accuracy.

Optimal Top-k Generation of Attribute Combinations Based on Ranked Lists

Jiaheng Lu, Renmin University of China; Pierre Senellart, Télécom ParisTech; Chunbin Lin, Renmin University of China; Xiaoyong Du, Renmin University of China; Shan Wang, Renmin University of China; Xinxing Chen, Renmin University of China

In this work, we study a novel query type, called top-k,m queries. Suppose we are given a set of groups and each group contains a set of attributes, each of which is associated with a ranked list of tuples, with ID and score. All lists are ranked in decreasing order of the scores of tuples. We are interested in finding the best combinations of attributes, each combination involving one attribute from each group. More specifically, we want the top-k combinations of attributes according to the corresponding top-m tuples with matching IDs. This problem has a wide range of applications from databases to search engines on traditional and non-traditional types of data (relational data, XML, text, etc.). We show that a straightforward extension of an optimal top-*k* algorithm, the Threshold Algorithm (TA), has shortcomings in solving the \topkm{} problem, as it needs to compute a large number of intermediate results for each combination and reads moreinputs than needed. To overcome this weakness, we provide here, for the first time, a provably instance-optimal algorithm and further develop

optimizations for efficient query evaluation to reduce computational and memory costs and the number of accesses. We demonstrate experimentally the scalability and efficiency of our algorithms over three real applications.

Top-k Bounded Diversification

Piero Fraternali, Politecnico di Milano; Davide Martinenghi, Politecnico di Milano; Marco Tagliasacchi, Politecnico di Milano

This paper investigates diversity queries over objects embedded in a low-dimensional vector space. An interesting case is provided by spatial Web objects, which are produced in great quantity by location-based services that let users attach content to places, and arise also in trip planning, news analysis, and real estate scenarios. The targeted queries aim at retrieving the best set of objects relevant to given user criteria and well distributed over a region of interest. Such queries are a particular case of diversified top-k queries, for which existing methods are too costly, as they evaluate diversity by accessing and scanning all relevant objects, even if only a small subset is needed. We therefore introduce Space Partitioning and Probing (SPP), an algorithm that minimizes the number of accessed objects while finding exactly the same result as MMR, the most popular diversification algorithm. SPP belongs to a family of algorithms that rely only on score-based and distance-based access methods, which are available in most geo-referenced Web data sources, and do not require retrieving all the relevant objects. Experiments show that SPP significantly reduces the number of accessed objects while incurring a very low computational overhead.

SIGMOD Research 13: Temporal and Graph Databases

Temporal Alignment

Anton Dignös, University of Zürich; Michael H. Böhlen, University of Zürich; Johann Gamper, Free University of Bozen-Bolzano

In order to process interval timestamped data, the sequenced semantics has been proposed. This paper presents a relational algebra solution that provides native support for the three properties of the sequenced semantics: snapshot reducibility, extended snapshot reducibility, and change preservation. We introduce two temporal primitives, temporal splitter and temporal aligner, and define rules that use these primitives to reduce the operators of a temporal algebra to their nontemporal counterparts. Our solution supports the three properties of the sequenced semantics through interval adjustment and timestamp propagation. We have implemented the temporal primitives and reduction rules in the kernel of PostgreSQL to get native database support for processing interval timestamped data. The support is comprehensive and includes outer joins, antijoins, and aggregations with predicates and functions over the time intervals of argument relations. The implementation and empirical evaluation confirms effectiveness and scalability of our solution that leverages existing database query optimization techniques.

A Highway-Centric Labeling Approach for Answering Distance Queries on Large Sparse Graphs

Ruoming Jin, Kent State University; Ning Ruan, Kent State University; Yang Xiang, Ohio State University; Victor Lee, Kent State University

The distance query, which asks the length of the shortest path from a vertex *u* to another vertex *v*, has applications ranging from link analysis, semantic web and other ontology processing, to social network operations. Here, we propose a novel labeling scheme, referred to as *Highway-Centric Labeling*, for answering distance queries in a large sparse graph.
It empowers the distance labeling with a highway structure and leverages a novel bipartite set cover framework/algorithm. Highway-centric labeling provides better labeling size than the state-of-the-art 2-hop labeling, theoretically and empirically. It also offers both exact distance and approximate distance with bounded accuracy. A detailed experimental evaluation on both synthetic and real datasets demonstrates that highway-centric labeling can outperform the state-of-the-art distance computation approaches in terms of both index size and query time.

Efficient Processing of Distance Queries in Large Graphs: A Vertex Cover Approach

James Cheng, Nanyang Technological University; Yiping Ke, Institute of High Performance Computing, Singapore; Shumo Chu, Nanyang Technological University; Carter Cheng, Nanyang Technological University

We propose a novel disk-based index for processing single-source shortest path or distance queries. The index is useful in a wide range of important applications (e.g., network analysis, routing planning, etc.). Our index is a tree-structured index constructed based on the concept of vertex cover. We propose an I/O-efficient algorithm to construct the index when the input graph is too large to fit in main memory. We give detailed analysis of I/O and CPU complexity for both index construction and query processing, and verify the efficiency of our index for query processing in massive real-world graphs.

SIGMOD Research 14: Information Retrieval and Text Mining

Aggregate Suppression for Enterprise Search Engines

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

Many enterprise websites provide search engines to facilitate customer access to their underlying documents or data. With the web interface of such a search engine, a customer can specify one or a few keywords that he/she is interested in; and the search engine returns a list of documents/tuples matching the user-specified keywords, sorted by an often-proprietary scoring function.
It was traditionally believed that, because of its highly-restrictive interface (i.e., keyword search only, no SQL-style queries), such a search engine serves its purpose of answering individual keyword-search queries without disclosing big-picture aggregates over the data which, as we shall show in the paper, may incur significant privacy concerns to the enterprise. Nonetheless, recent work on sampling and aggregate estimation over a search engine's corpus through its keyword-search interface transcends this traditional belief. In this paper, we consider a novel problem of suppressing sensitive aggregates for enterprise search engines while maintaining the quality of answers provided to individual keyword-search queries. We demonstrate the effectiveness and efficiency of our novel techniques through theoretical analysis and extensive experimental studies.

Probase: A Probabilistic Taxonomy for Text Understanding

Wentao Wu, University of Wisconsin-Madison; Hongsong Li, Microsoft Research Asia; Haixun Wang, Microsoft Research Asia; Kenny Q. Zhu, Shanghai Jiao Tong University

Knowledge is indispensable to understanding. The ongoing information explosion highlights the need to enable machines to better understand electronic text in human language. Much work has been devoted to creating universal ontologies or taxonomies for this purpose. However, none of the existing ontologies has the needed depth and breadth for universal understanding. In this paper, we present a universal, probabilistic taxonomy that is more comprehensive than any existing ones. It contains 2.7 million concepts harnessed automatically from a corpus of 1.68 billion web pages. Unlike traditional taxonomies that treat knowledge as black and white, it uses probabilities to model inconsistent, ambiguous and uncertain information it contains. We present details of how the taxonomy is constructed, its probabilistic modeling, and its potential applications in text understanding.

Optimizing Index for Taxonomy Keyword Search

Bolin Ding, University of Illinois at Urbana Champaign; Haixun Wang, Microsoft Research Asia; Ruoming Jin, Kent State University; Jiawei Han, University of Illinois at Urbana Champaign; Zhongyuan Wang, Microsoft Research Asia

Query substitution is an important problem in information retrieval. Much work focuses on how to find substitutes for any given query. In this paper, we study how to efficiently process a keyword query whose substitutes are defined by a given taxonomy. This problem is challenging because each term in a query can have a large number of substitutes, and the original query can be rewritten into any of their combinations. We propose to build an additional index (besides inverted index) to efficiently process queries. For a query workload, we formulate an optimization problem which chooses the additional index structure, aiming at minimizing the query evaluation cost, under given index space constraints. We show the NP-hardness of the problem, and propose a pseudo-polynomial time algorithm using dynamic programming, as well as an $\frac{1}{4}(1-1/e)$-approximation algorithm to solve the problem. Experimental results show that, with only 10% additional index space, our approach can greatly reduce the query evaluation cost.

SIGMOD Research 15: Social Networks and Graph Databases II

A Model-Based Approach to Attributed Graph Clustering

Zhiqiang Xu, Nanyang Technological University; Yiping Ke, Institute of High Performance Computing, Singapore; Yi Wang, National University of Singapore; Hong Cheng, The Chinese University of Hong Kong; James Cheng, Nanyang Technological University

Graph clustering, also known as community detection, is a long-standing problem in data mining. However, with the proliferation of rich attribute information available for objects in real-world graphs, how to leverage structural and attribute information for clustering attributed graphs becomes a new challenge. Most existing works take a distance-based approach. They proposed various distance measures to combine structural and attribute information. In this paper, we consider an alternative view and propose a model-based approach to attributed graph clustering. We develop a Bayesian probabilistic model for attributed graphs. The model provides a principled and natural framework for capturing both structural and attribute aspects of a graph, while avoiding the artificial design of a distance measure. Clustering with the proposed model can be transformed into a probabilistic inference problem, for which we devise an efficient variational algorithm.

Experimental results on large real-world datasets demonstrate that our method significantly outperforms the state-of-art distance-based attributed graph clustering method.

Towards Effective Partition Management for Large Graphs

Shengqi Yang, University of California, Santa Barbara; Xifeng Yan, University of California, Santa Barbara; Bo Zong, University of California, Santa Barbara; Arijit Khan, University of California, Santa Barbara

Searching and mining large graphs today is critical to a variety of application domains, ranging from community detection in social networks, to de novo genome sequence assembly. Scalable processing of large graphs requires careful partitioning and distribution of graphs across clusters. In this paper, we investigate the problem of managing large-scale graphs in clusters and study access characteristics of local graph queries such as breadth-first search, random walk, and SPARQL queries, which are popular in real applications. These queries exhibit strong access locality, and therefore require specific data partitioning strategies. In this work, we propose a Self Evolving Distributed Graph Management Environment (\Sedge), to minimize inter-machine communication during graph query processing in multiple machines. In order to improve query response time and throughput, \Sedge\ introduces a two-level partition management architecture with complimentary primary partitions and dynamic secondary partitions. These two kinds of partitions are able to adapt in real time to changes in query workload. \Sedge\ also includes a set of workload analyzing algorithms whose time complexity is linear or sublinear to graph size. Empirical results show that it significantly improves distributed graph processing on today's commodity clusters.

TreeSpan: Efficiently Computing Similarity All-Matching

Gaoping Zhu, University of New South Wales; Xuemin Lin, University of New South Wales; Ke Zhu, University of New South Wales; Wenjie Zhang, University of New South Wales; Jeffrey Xu Yu, The Chinese University of Hong Kong

Given a query graph *q* and a data graph *G*, computing all occurrences of *q* in *G*, namely exact all-matching, is fundamental in graph data analysis with a wide spectrum of real applications. It is challenging since even finding one occurrence of *q* in *G* (subgraph isomorphism test) is NP-Complete. Consider that in many real applications, exploratory queries from users are often inaccurate to express their real demands. In this paper, we study the problem of efficiently computing all approximate occurrences of *q* in *G*. Particularly, we study the problem of efficiently retrieving all matches of *q* in *G* with the number of possible missing edges bounded by a given threshold θ, namely similarity all-matching. The problem of similarity all-matching is harder than the problem of exact all-matching since it covers the problem of exact all-matching as a special case with $\theta = 0$.

In this paper, we develop a novel paradigm to conduct similarity all-matching.Specifically, we propose to use a minimal set *QT* of spanning trees in *q* to *cover* all connected subgraphs *q'* of *q* missing at most θ edges; that is, each *q'* is spanned by a spanning tree in *QT*. Then, we conduct exact all-matching for each spanning tree in QT to induce all similarity matches.A rigid theoretic analysis shows that our new search paradigm significantly reduces the times of conducting exact all-matching against the existing techniques. To further speed-up the computation, we develop new filtering, computation sharing, and search ordering techniques. Our comprehensive experiments on both real and synthetic datasets demonstrate that our techniques outperform the state of the art technique by 7 orders of magnitude.

SIGMOD Research 16: Indexing and Physical Database Design II

Locality-Sensitive Hashing Scheme Based on Dynamic Collision Counting

Junho Gan, Sun Yat-Sen University; Jianlin Feng, Sun Yat-Sen University; Qiong Fang, Hong Kong University of Science and Technology; Wilfred Ng, Hong Kong University of Science and Technology

Locality-Sensitive Hashing (LSH) and its variants are well-known methods for solving the *c*-approximate NN Search problem in high-dimensional space. Traditionally, several LSH functions are concatenated to form a "static" compound hash function for building a hash table. In this paper, we propose to use a base of *m* single LSH functions to construct "*dynamic*" compound hash functions, and define a new LSH scheme called *Collision Counting LSH* (C2LSH). If the number of LSH functions under which a data object *o* collides with a query object *q* is greater than a pre-specified collision threhold l, then *o* can be regarded as a good candidate of *c*-approximate NN of *q*. This is the basic idea of C2LSH.

Our theoretical studies show that, by appropriately choosing the size of LSH function base *m* and the collision threshold l, C2LSH can have a guarantee on query quality. Notably, the parameter *m* is not affected by dimensionality of data objects, which makes C2LSH especially good for high dimensional NN search. The experimental studies based on synthetic datasets and four real datasets have shown that C2LSH outperforms the state of the art method LSB-forest in high dimensional space.

Efficient External-Memory Bisimulation on DAGs

Jelle Hellings, Hasselt University and Transnational University of Limburg; George H. L. Fletcher, Eindhoven University of Technology; Herman Haverkort, Eindhoven University of Technology

In this paper, we introduce the first efficient external-memory algorithm to compute the bisimilarity equivalence classes of a directed acyclic graph (DAG). DAGs are commonly used to model data in a wide variety of practical applications, ranging from XML documents and data provenance models, to web taxonomies and scientific workflows. In the study of efficient reasoning over massive graphs, the notion of node bisimilarity plays a central role. For example, grouping together bisimilar nodes in an XML data set is the first step in many sophisticated approaches to building indexing data structures for efficient XPath query evaluation. To date, however, only internal-memory bisimulation algorithms have been investigated. As the size of real-world DAG data sets often exceeds available main memory, storage in external memory becomes necessary. Hence, there is a practical need for an efficient approach to computing bisimulation in external memory.
Our general algorithm has a worst-case IO-complexity of $O(\Sort(|N| + |E|))$, where $|N|$ and $|E|$ are the numbers of nodes and edges, resp., in the data graph and $\Sort(n)$ is the number of accesses to external memory needed to sort an input of size n. We also study specializations of this algorithm to common variations of bisimulation for tree-structured XML data sets. We empirically verify efficient performance of the algorithms on graphs and XML documents having billions of nodes and edges, and find that the algorithms can process such graphs efficiently even when very limited internal memory is available. The proposed algorithms are simple enough for practical implementation and use, and open the door for further study of external-memory bisimulation algorithms. To this end, the full open-source C++ implementation has been made freely available.

Materialized View Selection for XQuery Workloads

Asterios Katsifodimos, INRIA Saclay and Université Paris-Sud; Ioana Manolescu, INRIA Saclay and Université Paris-Sud; Vasilis Vassalos, Athens University of Economics and Business

The efficient processing of XQuery still poses significant challenges. A particularly effective technique to improve XQuery processing performance consists of using materialized views to answer queries. In this work, we consider the problem of choosing the best views to materialize within a given space budget in order to improve the performance of a query workload. The paper is the first to address the view selection problem for queries and views with value joins and multiple return nodes. The challenges we face stem from the expressive power and features of both the query and view languages and from the size of the search space of candidate views to materialize. While the general problem has prohibitive complexity, we propose and study a heuristic algorithm and demonstrate its superior performance compared to the state of the art.

PODS PAPER ABSTRACTS

PODS Session 1: Streaming
Session Chair: Milan Vojnovic (Microsoft Research, Cambridge)

Graph Sketches: Sparsification, Spanners, and Subgraphs

Kook Jin Ahn, University of Pennsylvania; Sudipto Guha, University of Pennsylvania; Andrew McGregor, University of Massachusetts, Amherst

When processing massive data sets, a core task is to construct *synopses* of the data. To be useful, a synopsis data structure should be easy to construct while also yielding good approximations of the relevant properties of the data set. A particularly useful class of synopses are *sketches*, i.e., those based on linear projections of the data. These are applicable in many models including various parallel, stream, and compressed sensing settings. A rich body of analytic and empirical work exists for sketching numerical data such as the frequencies of a set of entities. Our work investigates *graph sketching* where the graphs of interest encode the relationships between these entities. The main challenge is to capture this richer structure and build the necessary synopses with only linear measurements.

In this paper we consider properties of graphs including the size of the cuts, the distances between nodes, and the prevalence of dense sub-graphs. Our main result is a sketch-based sparsifier construction: we show that $O(n\varepsilon^{-2})$ random linear projections of a graph on n nodes suffice to $(1+\varepsilon)$ approximate *all* cut values. Similarly, we show that $O(\varepsilon^{-2})$ linear projections suffice for (additively) approximating the fraction of induced sub-graphs that match a given pattern such as a small clique. Finally, for distance estimation we present sketch-based spanner constructions. In this last result the sketches are adaptive, i.e., the linear projections are performed in a small number of batches where each projection may be chosen dependent on the outcome of earlier sketches. All of the above results immediately give rise to data stream algorithms that also apply to dynamic graph streams where edges are both inserted and deleted. The non-adaptive sketches, such as those for sparsification and subgraphs, give us single-pass algorithms for distributed data streams with insertion and deletions. The adaptive sketches can be used to analyze MapReduce algorithms that use a small number of rounds.

Approximating and Testing k-Histogram Distributions in Sub-linear Time

Piotr Indyk, MIT; Reut Levi, Tel Aviv University; Ronitt Rubinfeld, MIT

A discrete distribution p, over $[n]$, is a k-histogram if its probability distribution function can be represented as a piece-wise constant function with k pieces. Such a function is represented by a list of k intervals and k corresponding values. We consider the following problem: given a collection of samples from a distribution p, find a k-histogram that (approximately) minimizes the l_2 distance to the distribution p. We give time and sample efficient algorithms for this problem. We further provide algorithms that distinguish distributions that have the property of being a k-histogram from distributions that are ε-far from any k-histogram in the l_1 distance and l_2 distance respectively.

Mergeable Summaries

Pankaj Agarwal, Duke University; Graham Cormode, AT&T Labs-Research; Zengfeng Huang, The Hong Kong University of Science and Technology; Jeff Phillips, University of Utah; Zhewei Wei, The Hong Kong University of Science and Technology; Ke Yi, The Hong Kong University of Science and Technology

We study the *mergeability* of data summaries. Informally speaking, mergeability requires that, given two summaries on two data sets, there is a way to merge the two summaries into a single summary on the union of the two data sets, while preserving the error and size guarantees. This property means that the summaries can be merged in a way like other algebraic operators such as sum and max, which is especially useful for computing summaries on massive distributed data. Several data summaries are trivially mergeable by construction, most notably all the *sketches* that are linear functions of the data sets. But some other fundamental ones like those for heavy hitters and quantiles, are not (known to be) mergeable. In this paper, we demonstrate that these summaries are indeed mergeable or can be made mergeable after appropriate modifications. Specifically, we show that for ε-approximate heavy hitters, there is a deterministic mergeable summary of size $O(1/\varepsilon)$; for ε-approximate quantiles, there is a deterministic summary of size $O(1/\varepsilon \log(\varepsilon n))$ that has a restricted form of mergeability, and a randomized one of size $O(1/\varepsilon \log^{3/2} 1/\varepsilon)$ with full mergeability. We also extend our results to geometric summaries such as ε-approximations and ε-kernels. We also achieve two results of independent interest: (1) we provide the best known randomized streaming bound for ε-approximate quantiles that depends only on ε, of size $O(1/\varepsilon \log^{3/2} 1/\varepsilon)$, and (2) we demonstrate that the MG and the SpaceSaving summaries for heavy hitters are isomorphic.

PODS Session 2: Awards Session
Session Chair: Richard Hull (IBM T. J. Watson Research Center)

The ACM PODS Alberto O. Mendelzon Test-of-Time Award: Containment and Equivalence for an XPath Fragment

Gerome Miklau, University of Massachusetts, Amherst; Dan Suciu, University of Washington

The Alberto O. Mendelzon Test-of-Time Award Committee for 2012 has decided to select the above paper as the award winner for 2012. The paper studied static analysis problems for XPath, a query language at the core of processing XML documents and XML document databases. XPath, an important paradigm of a query language for semi-structured data, is designed with tree-navigation in mind and supports such navigation along three axes: ancestor-descendant, branching, and wildcards.

Miklau and Suciu established that if all three axes are allowed, then the query-containment problem for XPath queries is coNP-complete. Furthermore, this intractability persists even when certain tight bounds on the number of wildcards and the number of branches are imposed. These results shed light on the boundary between tractability and intractability for XPath query containment, since it was previously known that the containment problem was solvable in polynomial time for XPath queries in which any two of the three axes are allowed.

Both the paper in the PODS 2002 proceedings and its subsequent full version in the Journal of the Association for Computing Machinery have received hundreds of citations each. Moreover, this work initiated a fruitful line of research on the static analysis of XML query languages that brought together researchers from database theory and automata theory.

Best Paper Award: Worst-Case Optimal Join Algorithms

Hung Q. Ngo, University at Buffalo, SUNY; Ely Porat, Bar-Ilan University; Christopher Ré, University of Wisconsin-Madison; Atri Rudra, University at Buffalo, SUNY

Efficient join processing is one of the most fundamental and well-studied tasks in database research. In this work, we examine algorithms for natural join queries over many relations and describe a novel algorithm to process these queries optimally in terms of worst-case data complexity. Our result builds on recent work by Atserias, Grohe, and Marx, who gave bounds on the size of a full conjunctive query in terms of the sizes of the individual relations in the body of the query. These bounds, however, are not constructive: they rely on Shearer's entropy inequality which is information-theoretic. Thus, the previous results leave open the question of whether there exist algorithms whose running time achieve these optimal bounds. An answer to this question may be interesting to database practice, as we show in this paper that any project-join plan is polynomially slower than the optimal bound for some queries. We construct an algorithm whose running time is worst-case optimal for all natural join queries. Our result may be of independent interest, as our algorithm also yields a constructive proof of the general fractional cover bound by Atserias, Grohe, and Marx without using Shearer's inequality. In addition, we show that this bound is equivalent to a geometric inequality by Bollobs and Thomason, one of whose special cases is the famous Loomis-Whitney inequality. Hence, our results algorithmically prove these inequalities as well. Finally, we discuss how our algorithm can be used to compute a relaxed notion of joins.

Regular Paper: Deterministic Regular Expressions in Linear Time

Benoit Groz, INRIA and University of Lille; Sebastian Maneth, NICTA and UNSW; Slawek Staworko, INRIA and University of Lille

Deterministic regular expressions are widely used in XML processing. For instance, all regular expressions in DTDs and XML Schemas are required to be deterministic. In this paper we show that determinism of a regular expression e can be tested in linear time. The best known algorithms, based on the Glushkov automaton, require $O(\sigma|e|)$ time, where σ is the number of distinct symbols in e. We further show that matching a word w against an expression e can be achieved in combined linear time $O(|e|+|w|)$, for a wide range of deterministic regular expressions: (i) star-free (for multiple input words), (ii) bounded-occurrence, i.e., expressions in which each symbol appears a bounded number of times, and (iii) bounded plus-depth, i.e., expressions in which the nesting depth of alternating plus (union) and concatenation symbols is bounded. Our algorithms use a new structural decomposition of the parse tree of e. For matching arbitrary deterministic regular expressions we present an $O(|e| + |w| \log \log |e|)$ time algorithm.

PODS Session 3: Tutorial Session 1
Session Chair: Phokion Kolaitis (University of California, Santa Cruz and IBM Almaden Research Center)

Invited Tutorial: Linguistic Foundations for Bidirectional Transformations

Benjamin C. Pierce, University of Pennsylvania

Computing is full of situations where two different structures must be "connected" in such a way that updates to each can be propagated to the other. This is a generalization of the classical *view update problem*, which has been studied for decades in

the database community; more recently, related problems have attracted considerable interest in other areas, including programming languages, software model transformation, user interfaces and system configuration. Among the fruits of this cross-pollination has been the development of a *linguistic* perspective on the problem. Rather than taking some view definition language as fixed (e.g., choosing some subset of relational algebra) and looking for tractable ways of "inverting" view definitions to propagate updates from view to source, we can directly design new *bidirectional programming languages* in which every expression defines a *pair* of functions mapping updates on one structure to updates on the other. Such structures are often called *lenses*.

The foundational theory of lenses has been studied extensively, and lens-based language designs have been developed in several domains, including strings, trees, relations, graphs, and software models. These languages share some common elements with modern functional languages—in particular, they come with very expressive type systems. In other respects, they are rather novel and surprising. This tutorial surveys recent developments in the theory of lenses and the practice of bidirectional programming languages.

PODS Session 4: Privacy and Semantic Web
Session Chair: Pierre Senellart (Télécom ParisTech)

The Power of the Dinur-Nissim Algorithm: Breaking Privacy of Statistical and Graph Databases

Krzysztof Choromanski, Columbia University; Tal Malkin, Columbia University

A few years ago, Dinur and Nissim (PODS, 2003) proposed an algorithm for breaking database privacy when statistical queries are answered with a perturbation error of magnitude $o(\sqrt{n})$ for a database of size n. This negative result is very strong in the sense that it completely reconstructs $\Omega(n)$ data bits with an algorithm that is simple, uses random queries, and does not put any restriction on the perturbation other than its magnitude. Their algorithm works for a model where the database consists of *bits*, and the statistical queries asked by the adversary are *sum queries* for a subset of locations.
In this paper we extend the attack to work for much more general settings in terms of the type of statistical query allowed, the database domain, and the general tradeoff between perturbation and privacy. Specifically, we prove:

- For queries of the type $\sum_{i=1}^{n} f_i x_i$ where f_i are i.i.d. and with a finite third moment and positive variance (this includes as a special case the sum queries of Dinur-Nissim and several subsequent extensions), we prove that the quadratic relation between the perturbation and what the adversary can reconstruct holds even for smaller perturbations, and even for a larger data domain. If f_i is Gaussian, Poissonian, or bounded and of positive variance, this holds for arbitrary data domains and perturbation; for other f_i this holds as long as the domain is not too large and the perturbation is not too small.
- A positive result showing that for a sum query the negative result mentioned above is tight. Specifically, we build a distribution on bit databases and an answering algorithm such that any adversary who wants to recover a little more than the negative result above allows, will not succeed except with negligible probability.
- We consider a richer class of summation queries, focusing on databases representing graphs, where each entry is an edge, and the query is a structural function of a subgraph. We show an attack that recovers a big portion of the graph edges, as long as the graph and the function satisfy certain properties.

The attacking algorithms in both our negative results are straightforward extensions of the Dinur-Nissim attack, based on asking ϕ-weighted queries or queries choosing a subgraph uniformly at random. The novelty of our work is in the analysis, showing that this simple attack is much more powerful than was previously known, as well as pointing to possible limits of this approach and putting forth new application domains such as graph problems (which may occur in social networks, Internet graphs, etc). These results may find applications not only for breaking privacy, but also in the positive direction, for recovering complicated structure information using inaccurate estimates about its substructures.

A Rigorous and Customizable Framework for Privacy

Daniel Kifer, Penn State University; Ashwin Machanavajjhala, Yahoo! Research

In this paper we introduce a new and general privacy framework called Pufferfish. The Pufferfish framework can be used to create new privacy definitions that are customized to the needs of a given application. The goal of Pufferfish is to allow experts in an application domain, who frequently do not have expertise in privacy, to develop rigorous privacy definitions for their data sharing needs. In addition to this, the Pufferfish framework can also be used to study existing privacy definitions. We illustrate the benefits with several applications of this privacy framework: we use it to formalize and prove the statement that differential privacy assumes independence between records, we use it to define and study the notion of composition in a broader context than before, we show how to apply it to protect unbounded continuous attributes and aggregate information, and we show how to use it to rigorously account for prior data releases.

Static Analysis and Optimization of Semantic Web Queries

Andrés Letelier, PUC Chile; Jorge Pérez, Universidad de Chile; Reinhard Pichler, Technische Universität Wien; Sebastian Skritek, Technische Universität Wien

Static analysis is a fundamental task in query optimization. In this paper we study static analysis and optimization techniques for SPARQL, which is the standard language for querying Semantic Web data. Of particular interest for us is the optionality feature in SPARQL. It is crucial in Semantic Web data management, where data sources are inherently incomplete and the user is usually interested in partial answers to queries. This feature is one of the most complicated constructors in SPARQL and also the one that makes this language depart from classical query languages such as relational conjunctive queries. We focus on the class of well-designed SPARQL queries, which has been proposed in the literature as a fragment of the language with good properties regarding query evaluation. We first propose a tree representation for SPARQL queries, called pattern trees, which captures the class of well-designed SPARQL graph patterns and which can be considered as a query execution plan. Among other results, we propose several transformation rules for pattern trees, a simple normal form, and study equivalence and containment. We also study the enumeration and counting problems for this class of queries.

The Complexity of Evaluating Path Expressions in SPARQL

Katja Losemann, Universität Bayreuth; Wim Martens, Universität Bayreuth

The World Wide Web Consortium (W3C) recently introduced property paths in SPARQL 1.1, a query language for RDF data. Property paths allow SPARQL queries to evaluate regular expressions over graph data. However, they differ from standard regular expressions in several notable aspects. For example, they have a limited form of negation, they have numerical occurrence indicators as syntactic sugar, and their semantics on graphs is defined in a non-standard manner. We formalize the W3C semantics of property paths and investigate various query evaluation problems on graphs. More specifically, let x and y be two nodes in an edge-labeled graph and r be an expression. We study the complexities of (1) deciding whether there exists a path from x to y that matches r and (2) counting how many paths from x to y match r. Our main results show that, compared to an alternative semantics of regular expressions on graphs, the complexity of (1) and (2) under W3C semantics is significantly higher. Whereas the alternative semantics remains in polynomial time for large fragments of expressions, the W3C semantics makes problems (1) and (2) intractable almost immediately. As a side-result, we prove that the membership problem for regular expressions with numerical occurrence indicators and negation is in polynomial time.

PODS Session 5: Range Queries
Session Chair: Srikanta Tirthapura (Iowa State)

Space-Efficient Range Reporting for Categorical Data

Yakov Nekrich, Universidad de Chile

In the colored (or categorical) range reporting problem the set of input points is partitioned into categories and stored in a data structure; a query asks for categories of points that belong to the query range. In this paper we study two-dimensional colored range reporting in the external memory model and present I/O-efficient data structures for this problem. In particular, we describe data structures that answer three-sided colored reporting queries in $O(K/B)$ I/Os and two-dimensional colored reporting queries in $O(\log_2 \log_B N + K/B)$ I/Os when points lie on an $N \times N$ grid, K is the number of reported colors, and B is the block size. The space usage of both data structures is close to optimal.

Dynamic Top-K Range Reporting in External Memory

Cheng Sheng, Chinese University of Hong Kong; Yufei Tao, Korea Advanced Institute of Science and Technology

In the *top-K range reporting* problem, the dataset contains N points in the real domain R, each of which is associated with a real-valued *score*. Given an interval $[x_1,x_2]$ in R and an integer $K \leq N$, a query returns the K points in $[x_1,x_2]$ having the smallest scores. We want to store the dataset in a structure so that queries can be answered efficiently. In the external memory model, the state of the art is a static structure that consumes $O(N/B)$ space, answers a query in $O(\log_B N + K/B)$ time, and can be constructed in $O(N + (N \log N / B) \log_{M/B} (N/B))$ time, where B is the size of a disk block, and M the size of memory. We present a fully-dynamic structure that retains the same space and query bounds, and can be updated in $O(\log^2_B N)$ amortized time per insertion and deletion. Our structure can be constructed in $O((N/B) \log_{M/B} (N/B))$ time.

Indexability of 2D Range Search Revisited: Constant Redundancy and Weak Indivisibility

Yufei Tao, Chinese University of Hong Kong

In the 2D *orthogonal range search* problem, we want to preprocess a set of 2D points so that, given any axis-parallel query rectangle, we can report all the data points in the rectangle efficiently. This paper presents a lower bound on the query time that can be achieved by any external memory structure that stores a point at most r times, where r is a constant integer.

Previous research has resolved the bound at two extremes: $r = 1$, and r being arbitrarily large. We, on the other hand, derive the explicit tradeoff at every specific r. A premise that lingers in existing studies is the so-called *indivisibility assumption*: all the information bits of a point are treated as an atom, i.e., they are always stored together in the same block. We partially remove this assumption by allowing a data structure to freely divide a point into individual bits stored in different blocks. The only assumption is that, those bits must be *retrieved* for reporting, as opposed to being *computed* – we refer to this requirement as the *weak indivisibility assumption*. We also describe structures to show that our lower bound is tight up to only a small factor.

PODS Session 6:
Session Chair: David Woodruff (IBM Almaden Research Center)

Tutorial 2: Approximate Computation and Implicit Regularization for Very Large-scale Data Analysis

Michael W. Mahoney, Stanford University

Database theory and database practice are typically the domain of computer scientists who adopt what may be termed an algorithmic perspective on their data. This perspective is very different than the more statistical perspective adopted by statisticians, scientific computers, machine learners, and other who work on what may be broadly termed statistical data analysis. In this article, I will address fundamental aspects of this algorithmic-statistical disconnect, with an eye to bridging the gap between these two very different approaches. A concept that lies at the heart of this disconnect is that of statistical regularization, a notion that has to do with how robust is the output of an algorithm to the noise properties of the input data. Although it is nearly completely absent from computer science, which historically has taken the input data as given and modeled algorithms discretely, regularization in one form or another is central to nearly every application domain that applies algorithms to noisy data. By using several case studies, I will illustrate, both theoretically and empirically, the nonobvious fact that approximate computation, in and of itself, can implicitly lead to statistical regularization. This and other recent work suggests that, by exploiting in a more principled way the statistical properties implicit in worst-case algorithms, one can in many cases satisfy the bicriteria of having algorithms that are scalable to very large-scale databases and that also have good inferential or predictive properties.

Regular Paper: Max-Sum Diversification, Monotone Submodular Functions and Dynamic Updates

Allan Borodin, University of Toronto; Hyun Chul Lee, LinkedIn Corporation; Yuli Ye, University of Toronto

Result diversification has many important applications in databases, operations research, information retrieval, and finance. In this paper, we study and extend a particular version of result diversification, known as max-sum diversification. More specifically, we consider the setting where we are given a set of elements in a metric space and a set valuation function f defined on every subset. For any given subset S, the overall objective is a linear combination of $f(S)$ and the sum of the distances induced by S. The goal is to find a subset S satisfying some constraints that maximizes the overall objective. This problem is first studied by Gollapudi and Sharma for modular set functions and for sets satisfying a cardinality constraint (uniform matroids). In their paper, they give a 2-approximation algorithm by reducing to an earlier result by Hassin, Rubinstein and Tamir. The first part of this paper considers an extension of the modular case to the monotone submodular case, for which the algorithm by Gollapudi and Sharma no longer applies. Interestingly, we are able to maintain the same 2-approximation using a natural, but different greedy algorithm. We then further extend the problem by considering any matroid constraint and show that a natural single swap local search algorithm provides a 2-approximation in this more general setting. This extends the Nemhauser, Wolsey and Fisher approximation result for the problem of submodular function maximization subject to a matroid constraint (without the distance function component). The second part of the paper focuses on dynamic updates for the modular case. Suppose we have a good initial approximate solution and then there is a single weight-perturbation either on the valuation of an element or on the distance between two elements. Given that users expect some stability in the results they see, we ask how easy is it to maintain a good approximation without significantly changing the initial set. We measure this by the number of updates, where each update is a swap of a single element in the current solution with a single element outside the current solution. We show that we can maintain an approximation ratio of 3 by just a single update if the perturbation is not too large.

PODS Session 7: Views and Data Exchange
Session Chair: Jan Paredaens (University of Antwerp)

Query-Based Data Pricing

Paraschos Koutris, University of Washington; Prasang Upadhyaya, University of Washington; Magdalena Balazinska, University of Washington; Bill Howe, University of Washington; Dan Suciu, University of Washington

Data is increasingly being bought and sold online, and Web-based marketplace services have emerged to facilitate these activities. However, current mechanisms for pricing data are very simple: buyers can choose only from a set of explicit views, each with a specific price. In this paper, we propose a framework for pricing data on the Internet that, given the price

of a few views, allows the price of any query to be derived automatically. We call this capability "query-based pricing." We first identify two important properties that the pricing function must satisfy, called *arbitrage-free* and *discount-free*. Then, we prove that there exists a unique function that satisfies these properties and extends the seller's explicit prices to all queries. When both the views and the query are Unions of Conjunctive Queries, the complexity of computing the price is high. To ensure tractability, we restrict the explicit prices to be defined only on selection views (which is the common practice today). We give an algorithm with polynomial time data complexity for computing the price of any chain query by reducing the problem to network flow. Furthermore, we completely characterize the class of Conjunctive Queries without self-joins that have PTIME data complexity (this class is slightly larger than chain queries), and prove that pricing all other queries is NP-complete, thus establishing a dichotomy on the complexity of the pricing problem when all views are selection queries.

Local Transformations and Conjunctive-Query Equivalence

Ronald Fagin, IBM Almaden Research Center; Phokion G. Kolaitis, University of California, Santa Cruz and IBM Almaden Research Center

Over the past several decades, the study of conjunctive queries has occupied a central place in the theory and practice of database systems. In recent years, conjunctive queries have played a prominent role in the design and use of schema mappings for data integration and data exchange tasks. In this paper, we investigate several different aspects of conjunctive-query equivalence in the context of schema mappings and data exchange.

In the first part of the paper, we introduce and study a notion of a local transformation between database instances that is based on conjunctive-query equivalence. We show that the chase procedure for GLAV mappings (that is, schema mappings specified by source-to-target tuple-generating dependencies) is a local transformation with respect to conjunctive-query equivalence. This means that the chase procedure preserves bounded conjunctive-query equivalence, that is, if two source instances are indistinguishable using conjunctive queries of a sufficiently large size, then the target instances obtained by chasing these two source instances are also indistinguishable using conjunctive queries of a given size. Moreover, we obtain polynomial bounds on the level of indistinguishability between source instances needed to guarantee indistinguishability between the target instances produced by the chase. The locality of the chase extends to schema mappings specified by a second-order tuple-generating dependency (SO tgd), but does not hold for schema mappings whose specification includes target constraints.

In the second part of the paper, we take a closer look at the composition of two GLAV mappings. In particular, we break GLAV mappings into a small number of well-studied classes (including LAV and GAV), and complete the picture as to when the composition of schema mappings from these various classes can be guaranteed to be a GLAV mapping, and when they can be guaranteed to be conjunctive-query equivalent to a GLAV mapping.

We also show that the following problem is decidable: given a schema mapping specified by an SO tgd and a GLAV mapping, are they conjunctive-query equivalent? In contrast, the following problem is known to be undecidable: given a schema mapping specified by an SO tgd and a GLAV mapping, are they logically equivalent?

A Dichotomy in the Complexity of Deletion Propagation with Functional Dependencies

Benny Kimelfeld, IBM Almaden Research Center

A classical variant of the view-update problem is deletion propagation, where tuples from the database are deleted in order to realize a desired deletion of a tuple from the view. This operation may cause a (sometimes necessary) side effect—deletion of additional tuples from the view, besides the intentionally deleted one. The goal is to propagate deletion so as to maximize the number of tuples that remain in the view. In this paper, a view is defined by a self-join-free conjunctive query (sjf-CQ) over a schema with functional dependencies. A condition is formulated on the schema and view definition at hand, and the following dichotomy in complexity is established. If the condition is met, then deletion propagation is solvable in polynomial time by an extremely simple algorithm (very similar to the one observed by Buneman et al.). If the condition is violated, then the problem is NP-hard, and it is even hard to realize an approximation ratio that is better than some constant; moreover, deciding whether there is a side-effect-free solution is NP-complete. This result generalizes a recent result by Kimelfeld et al., who ignore functional dependencies. For the class of sjf-CQs, it also generalizes a result by Cong et al., stating that deletion propagation is in polynomial time if keys are preserved by the view.

PODS Session 8: Indexing
Session Chair: Yakov Nekrich (University of Bonn)

The Wavelet Trie: Maintaining an Indexed Sequence of Strings in Compressed Space

Roberto Grossi, Università di Pisa; Giuseppe Ottaviano, Università di Pisa

An indexed sequence of strings is a data structure for storing a string sequence that supports random access, searching, range counting and analytics operations, both for exact matches and prefix search. String sequences lie at the core of

column-oriented databases, log processing, and other storage and query tasks. In these applications each string can appear several times and the order of the strings in the sequence is relevant. The prefix structure of the strings is relevant as well: common prefixes are sought in strings to extract interesting features from the sequence. Moreover, space-efficiency is highly desirable as it translates directly into higher performance, since more data can fit in fast memory. We introduce and study the problem of compressed indexed sequence of strings, representing indexed sequences of strings in nearly-optimal compressed space, both in the static and dynamic settings, while preserving provably good performance for the supported operations. We present a new data structure for this problem, the Wavelet Trie, which combines the classical Patricia Trie with the Wavelet Tree, a succinct data structure for storing a compressed sequence. The resulting Wavelet Trie smoothly adapts to a sequence of strings that changes over time. It improves on the state-of-the-art compressed data structures by supporting a dynamic alphabet (i.e. the set of distinct strings) and prefix queries, both crucial requirements in the aforementioned applications, and on traditional indexes by reducing space occupancy to close to the entropy of the sequence.

On the Optimality of Clustering Properties of Space Filling Curves

Pan Xu, Iowa State University; Srikanta Tirthapura, Iowa State University

Space filling curves have for long been used in the design of data structures for multidimensional data. A fundamental quality metric of a space filling curve is its "clustering number" with respect to a class of queries, which is the average number of contiguous segments on the space filling curve that a query region can be partitioned into. We present a characterization of the clustering number of a general class of space filling curves, as well as the first non-trivial lower bounds on the clustering number for any space filling curve. Our results also answer an open problem that was posed by Jagadish in 1997.

Nearest-Neighbor Searching Under Uncertainty

Pankaj K. Agarwal, Duke University; Alon Efrat, The University of Arizona; Swaminathan Sankararaman, Duke University; Wuzhou Zhang, Duke University

Nearest-neighbor queries, which ask for returning the nearest neighbor of a query point in a set of points, are important and widely studied in many fields because of a wide range of applications. In many of these applications, such as sensor databases, location based services, face recognition, and mobile data, the location of data is imprecise. We therefore study nearest neighbor queries in a probabilistic framework in which the location of each input point and/or query point is specified as a probability density function and the goal is to return the point that minimizes the expected distance, which we refer to as the expected nearest neighbor (ENN). We present methods for computing an exact ENN or an ε-approximate ENN, for a given error parameter $0 < \varepsilon < 1$, under different distance functions. These methods build an index of near-linear size and answer ENN queries in polylogarithmic or sublinear time, depending on the underlying function. As far as we know, these are the first nontrivial methods for answering exact or ε-approximate ENN queries with provable performance guarantees.

PODS Session 9: Query Languages
Session Chair: Daniel Kifer (Penn State University)

Classification of Annotation Semirings over Query Containment

Egor V. Kostylev, University of Edinburgh; Juan L. Reutter, University of Edinburgh; András Z. Salamon, University of Edinburgh

We study the problem of query containment of (unions of) conjunctive queries over annotated databases. Annotations are typically attached to tuples and represent metadata such as probability, multiplicity, comments, or provenance. It is usually assumed that annotations are drawn from a commutative semiring. Such databases pose new challenges in query optimization, since many related fundamental tasks, such as query containment, have to be reconsidered in the presence of propagation of annotations. We axiomatize several classes of semirings for each of which containment of conjunctive queries is equivalent to existence of a particular type of homomorphism. For each of these types we also specify all semirings for which existence of a corresponding homomorphism is a sufficient (or necessary) condition for the containment. We exploit these techniques to develop new decision procedures for containment of unions of conjunctive queries and axiomatize corresponding classes of semirings. This generalizes previous approaches and allows us to improve known complexity bounds.

Efficient Approximations of Conjunctive Queries

Pablo Barceló, Universidad de Chile; Leonid Libkin, University of Edinburgh; Miguel Romero, Universidad de Chile

When finding exact answers to a query over a large database is infeasible, it is natural to approximate the query by a more efficient one that comes from a class with good bounds on the complexity of query evaluation. In this paper we study such

approximations for conjunctive queries. These queries are of special importance in databases, and we have a very good understanding of the classes that admit fast query evaluation, such as acyclic, or bounded (hyper)treewidth queries. We define approximations of a given query Q as queries from one of those classes that disagree with Q as little as possible. We mostly concentrate on approximations that are guaranteed to return correct answers. We prove that for the above classes of tractable conjunctive queries, approximations always exist, and are at most polynomial in the size of the original query. This follows from general results we establish that relate closure properties of classes of conjunctive queries to the existence of approximations. We also show that in many cases, the size of approximations is bounded by the size of the query they approximate. We establish a number of results showing how combinatorial properties of queries affect properties of their approximations, study bounds on the number of approximations, as well as the complexity of finding and identifying approximations. We also look at approximations that return all correct answers and study their properties.

On the Complexity of Package Recommendation Problems

Ting Deng, Beihang University; Wenfei Fan, University of Edinburgh; Floris Geerts, University of Antwerp

Recommendation systems aim to recommend items that are likely to be of interest to users. This paper investigates several issues fundamental to such systems. (1) We model recommendation systems for packages of items. We use queries to specify multi-criteria for item selections and express compatibility constraints on items in a package, and use functions to compute the cost and usefulness of items to a user. (2) We study recommendations of points of interest, to suggest top-k packages. We also investigate recommendations of top-k items, as a special case. In addition, when sensible suggestions cannot be found, we propose query relaxation recommendations to help users revise their selection criteria, or adjustment recommendations to guide vendors to modify their item collections. (3) We identify several problems, to decide whether a set of packages makes a top-k recommendation, whether a rating bound is maximum for selecting top-k packages, whether we can relax the selection query to find packages that users want, and whether we can update a bounded number of items such that the users' requirements can be satisfied. We also study function problems for computing top-k packages, and counting problems to find how many packages meet the user's criteria. (4) We establish the upper and lower bounds of these problems, all matching, for combined and data complexity. These results reveal the impact of variable sizes of packages, the presence of compatibility constraints, as well as a variety of query languages for specifying selection criteria and compatibility constraints, on the analyses of these problems.

PODS Session 10: Streaming and Aggregation
Session Chair: Yufei Tao (Chinese University of Hong Kong)

Space-Efficient Estimation of Statistics over Sub-Sampled Streams

Andrew McGregor, University of Massachusetts, Amherst; A. Pavan, Iowa State University; Srikanta Tirthapura, Iowa State University; David Woodruff, IBM Almaden Research Center

In many stream monitoring situations, the data arrival rate is so high that it is not even possible to observe each element of the stream. The most common solution is to sample a small fraction of the data stream and use the sample to infer properties and estimate aggregates of the original stream. However, the quantities that need to be computed on the sampled stream are often different from the original quantities of interest and their estimation requires new algorithms. We present upper and lower bounds (often matching) for estimating frequency moments, support size, entropy, and heavy hitters of the original stream from the data observed in the sampled stream.

Rectangle-Efficient Aggregation in Spatial Data Streams

Srikanta Tirthapura, Iowa State University; David Woodruff, IBM Almaden Research Center

We consider the estimation of aggregates over a data stream of multidimensional axis-aligned rectangles. Rectangles are a basic primitive object in spatial databases, and efficient aggregation of rectangles is a fundamental task. The data stream model has emerged as a de facto model for processing massive databases in which the data resides in external memory or the cloud and is streamed through main memory. For a point p, let $n(p)$ denote the sum of the weights of all rectangles in the stream that contain p. We give near-optimal solutions for basic problems, including (1) the k-th frequency moment $F_k = \Sigma_{\text{points } p} |n(p)|^k$, (2) the counting version of stabbing queries, which seeks an estimate of $n(p)$ given p, and (3) identification of heavy-hitters, i.e., points p for which $n(p)$ is large. An important special case of F_k is F_0, which corresponds to the volume of the union of the rectangles. This is a celebrated problem in computational geometry known as "Klee's measure problem", and our work yields the first solution in the streaming model for dimensions greater than one.

Randomized Algorithms for Tracking Distributed Count, Frequencies, and Ranks

Zengfeng Huang, Hong Kong University of Science and Technology; Ke Yi, Hong Kong University of Science and Technology; Qin Zhang, Aarhus University

We show that randomization can lead to significant improvements for a few fundamental problems in distributed tracking. Our basis is the *count-tracking* problem, where there are k players, each holding a counter n_i that gets incremented over

time, and the goal is to track an ε-approximation of their sum $n = \Sigma_i n_i$ continuously at all times, using minimum communication. While the deterministic communication complexity of the problem is $\Theta(k/\varepsilon \cdot \log N)$, where N is the final value of n when the tracking finishes, we show that with randomization, the communication cost can be reduced to $\Theta(\sqrt{k}/\varepsilon \cdot \log N)$. Our algorithm is simple and uses only O(1) space at each player, while the lower bound holds even assuming each player has infinite computing power. Then, we extend our techniques to two related distributed tracking problems: *frequency-tracking* and *rank-tracking*, and obtain similar improvements over previous deterministic algorithms. Both problems are of central importance in large data monitoring and analysis, and have been extensively studied in the literature.

Continuous Distributed Counting for Non-monotonic Streams

Zhenming Liu, Harvard University; Bozidar Radunovic, Microsoft Research; Milan Vojnovic, Microsoft Research

We consider the continual count tracking problem in a distributed environment where the input is an aggregate stream that originates from k distinct sites and the updates are allowed to be non-monotonic, i.e. both increments and decrements are allowed. The goal is to continually track the count within a prescribed relative accuracy ε at the lowest possible communication cost. Specifically, we consider an adversarial setting where the input values are selected and assigned to sites by an adversary but the order is according to a random permutation or is a random i.i.d process. The input stream of values is allowed to be non-monotonic with an unknown drift $-1 \le \mu \le 1$ where the case $\mu = 1$ corresponds to the special case of a monotonic stream of only non-negative updates. We show that a randomized algorithm guarantees to track the count accurately with high probability and has the expected communication cost $\sim O(\min\{\sqrt{k}/(|\mu|\ \varepsilon),\ \sqrt{k \times n}/\varepsilon,\ n\})$, for an input stream of length n, and establish matching lower bounds. This improves upon previously best known algorithm whose expected communication cost is $\sim\Theta(\min\{\sqrt{k}/\varepsilon, n\})$ that applies only to an important but more restrictive class of monotonic input streams, and our results are substantially more positive than the communication complexity of $\Omega(n)$ under fully adversarial input. We also show how our framework can also accommodate other types of random input streams, including fractional Brownian motion that has been widely used to model temporal long-range dependencies observed in many natural phenomena. Last but not least, we show how our non-monotonic counter can be applied to track the second frequency moment and to a Bayesian linear regression problem.

SIGMOD INDUSTRIAL PAPER ABSTRACTS

SIGMOD Industry 1: Databases in the Cloud

Amazon DynamoDB: A Seamlessly Scalable Non-Relational Datastore

Swami Sivasubramanian, Amazon

Reliability and scalability of an application is dependent on how its application state is managed. To run applications at massive scale requires one to operate datastores that can scale to operate seamlessly across thousands of servers and can deal with various failure modes such as server failures, datacenter failures and network partitions. The goal of Amazon DynamoDB is to eliminate this complexity and operational overhead for our customers by offering a seamlessly scalable database service. In this talk, I will talk about how developers can build applications on DynamoDB without having to deal with the complexity of operating a large scale database.

Efficient Transaction Processing in SAP HANA Database--The End of a Column Store Myth

Vishal Sikka, SAP; Franz Färber, SAP; Wolfgang Lehner, TUD/SAP; Sang Kyun Cha, SAP; Thomas Peh, SAP; Christof Bornhövd, SAP

The SAP HANA database is the core of SAP's new data management platform. The overall goal of the SAP HANA database is to provide a generic but powerful system for different query scenarios, both transactional and analytical, on the same data representation within a highly scalable execution environment. Within this paper, we highlight the main features that differentiate the SAP HANA database from classical relational database engines. Therefore, we outline the general architecture and design criteria of the SAP HANA in a first step. In a second step, we challenge the common belief that column store data structures are only superior in analytical workloads and not well suited for transactional workloads. We outline the concept of record life cycle management to use different storage formats for the different stages of a record. We not only discuss the general concept but also dive into some of the details of how to efficiently propagate records through their life cycle and moving database entries from write-optimized to read-optimized storage formats. In summary, the paper aims at illustrating how the SAP HANA database is able to efficiently work in analytical as well as transactional workload environments.

Walnut: A Unified Cloud Object Store

Jianjun Chen, Yahoo!; Chris Douglas, Yahoo!; Michi Mutsuzaki, Yahoo!; Patrick Quaid, Yahoo!; Raghu Ramakrishnan, Yahoo!; Sriram Rao, Yahoo!; Russell Sears, Yahoo!

Walnut is an object-store being developed at Yahoo! with the goal of serving as a common low-level storage layer for a variety of cloud data management systems including Hadoop (a MapReduce system), MObStor (a multimedia serving system), and PNUTS (an extended key-value serving system). Thus, a key performance challenge is to meet the latency and throughput requirements of the wide range of workloads commonly observed across these diverse systems. The motivation for Walnut is to leverage a carefully optimized low-level storage system, with support for elasticity and high-availability, across all of Yahoo!'s data clouds. This would enable sharing of hardware resources across hitherto siloed clouds of different types, offering greater potential for intelligent load balancing and efficient elastic operation, and simplify the operational tasks related to data storage.

In this paper, we discuss the motivation for unifying different storage clouds, describe the requirements of a common storage layer, and present the Walnut design, which uses a quorum-based replication protocol and one-hop direct client access to the data in most regular operations. A unique contribution of Walnut is its hybrid object strategy, which efficiently supports both small and large objects. We present experiments based on both synthetic and real data traces, showing that Walnut works well over a wide range of workloads, and can indeed serve as a common low-level storage layer across a range of cloud systems.

SIGMOD Industry 2: Social Media and Crowdsourcing

The Value of Social Media Data in Enterprise Applications

Shivakumar Vaithyanathan, IBM Almaden Research Center

Social media is an interactive vehicle for communication accessed on a daily basis by hundreds of millions of people. Unlike conventional media, which is a one-way street for information exchange, social media enables people to write content as well as provide feedback and recommend content to other users. There are multiple enterprise applications, such as customer retention, new customer acquisition, campaign management and lead generation that can significantly benefit from the consumer insights hidden in the massive amounts of social media content. Defining, extracting and representing entities such as people, organization and products, and their inter-relationships enables the building of comprehensive consumer profiles that can be leveraged in enterprise applications. Building these social media profiles requires a combination of text and entity analytics, while the utilization of such profiles makes heavy use of statistical models and

machine learning. In this talk I will briefly describe the work in progress at IBM Research - Almaden on how such consumer insights, both at the level of an individual and at the level of appropriate micro-segments, can be used in enterprise applications in companies ranging from movie studios to financial services and insurance companies. I will also provide a brief overview of text, entity and statistical modeling tools that can operate in a distributed fashion over very large amounts of data.

Anatomy of a Gift Recommendation Engine Powered by Social Media

Yannis Pavlidis, @WalmartLabs; Madhusudan Mathihalli, @WalmartLabs; Indrani Chakravarty, @WalmartLabs; Arvind Batra, @WalmartLabs; Ron Benson, @WalmartLabs; Ravi Raj, @WalmartLabs; Robert Yau, @WalmartLabs; Mike McKiernan, @WalmartLabs; Venky Harinarayan, @WalmartLabs; Anand Rajaraman, @WalmartLabs

More and more people conduct their shopping online, especially during the holiday season. Shopping online offers a lot of convenience, including the luxury of shopping from home, the ease of research, better prices, and in many cases access to unique products not available in stores.

One of the facets of shopping is gifting. Gifting may be the act of giving a present to somebody because of an event (e.g., birthday) or occasion (e.g., house warming party). People may also treat themselves or loved ones to a gift. Regardless of the occasion or the reason for gifting, there is often one common denominator: delight the receiver. The pursuit of delight can cause a great deal of stress and also be extremely time consuming as many people today either already have everything, or have easy access to everything.

The @WalmartLabs Gift Recommendation Engine and its first application, Shopycat, which is a gift finder application on Facebook, aim to find the right and wow gifts much easier and quicker than ever before, by taking into account social media interactions. In this paper we will begin by describing the Shopycat Social Gift Finder Facebook application. Next, we describe the components of the engine. Finally, we discuss the metrics used to evaluate the engine.

Building such a gift recommendation engine raises many challenges, in inferring user interests, computing the giftability of a product and an interest, and processing the big and fast data associated with social media. We briefly discuss our solutions to these challenges. Overall, our gift recommendation engine is an example that illustrates social commerce, a powerful emerging trend in e-commerce, and a major focus of @WalmartLabs.

Designing a Scalable Crowdsourcing Platform

Chris Van Pelt, CrowdFlower; Alex Sorokin, CrowdFlower

Computers are extremely efficient at crawling, storing and processing huge volumes of structured data. They are great at exploiting link structures to generate valuable knowledge. Yet there are plenty of data processing tasks that are difficult today. Labeling sentiment, moderating images, and mining structured content from the web are still too hard for computers. Automated techniques can get us a long way in some of those, but human inteligence is required when an accurate decision is ultimately important. In many cases that decision is easy for people and can be made quickly - in a few seconds to few minutes. By creating millions of simple online tasks we create a distributed computing machine. By shipping the tasks to millions of contributers around the globe, we make this human computer available 24/7 to make important decisions about your data. In this talk, I will describe our approach to designing CrowdFlower - a scalable crowdsourcing platform - as it evolved over the last 4 years. We think about crowdsourcing in terms of Quality, Cost and Speed. They are the ultimate design objectives of a human computer. Unfortunately, we cant have all 3. A general price-constrained task requiring 99.9% accuracy and 10 minute turnaround is not possible today. I will discuss design decisions behind CrowdFlower that allow us to pursue any two of these objectives. I will briefly present examples of common crowdsourced tasks and tools built into the platform to make the design of complex tasks easy, tools such as CrowdFlower Markup Language(CML). Quality control is the single most important challenge in Crowdsourcing. To enable an unidentified crowd of people to produce meaningful work, we must be certain that we can filter out bad contributors and produce high quality output. Initially we only used consensus. As the diversity and size of our crowd grew, so did the number of people attempting fraud. CrowdFlower developed Gold standard to block attempts of fraud. The use of gold allowed us to train contributors for the details of specific domains. By defining expected responses for a subset of the work and providing explanations of why a given response was expected, we are able distribute tasks to an ever-expanding anonymous workforce without sacrificing quality.

SIGMOD Industry 3: Modern RDBMSs

Query Optimization in Microsoft SQL Server PDW

Srinath Shankar, Microsoft; Rimma Nehme, Microsoft; Josep Aguilar-Saborit, Microsoft; Andrew Chung, Microsoft; Mostafa Elhemali, Microsoft; Alan Halverson, Microsoft; Eric Robinson, Microsoft; Mahadevan Sankara Subramanian, Microsoft; David DeWitt, Microsoft; César Galindo-Legaria, Microsoft

In recent years, Massively Parallel Processors have increasingly been used to manage and query vast amounts of data. Dramatic performance improvements are achieved through distributed execution of queries across many nodes. Query optimization for such system is a challenging and important problem.
In this paper we describe the Query Optimizer inside the SQL Server Parallel Data Warehouse product (PDW QO). We leverage existing QO technology in Microsoft SQL Server to implement a cost-based optimizer for distributed query execution. By properly abstracting metadata we can readily reuse existing logic for query simplification, space exploration and cardinality estimation. Unlike earlier approaches that simply parallelize the best serial plan, our optimizer considers a rich space of execution alternatives, and picks one based on a cost-model for the distributed execution environment. The result is a high-quality, effective query optimizer fordistributed query processing in an MPP.

F1—The Fault-Tolerant Distributed RDBMS Supporting Google's Ad Business

Jeff Shute, Google; Mircea Oancea, Google; Stephan Ellner, Google; Ben Handy, Google; Eric Rollins, Google; Bart Samwel, Google; Radek Vingralek, Google; Chad Whipkey, Google; Xin Chen, Google; Beat Jegerlehner, Google; Kyle Littlefield, Google; Phoenix Tong, Google

Many of the services that are critical to Google's ad business have historically been backed by MySQL. We have recently migrated several of these services to F1, a new RDBMS developed at Google. F1 implements rich relational database features, including a strictly enforced schema, a powerful parallel SQL query engine, general transactions, change tracking and notification, and indexing, and is built on top of a highly distributed storage system that scales on standard hardware in Google data centers. The store is dynamically sharded, supports transactionally-consistent replication across data centers, and is able to handle data center outages without data loss.
The strong consistency properties of F1 and its storage system come at the cost of higher write latencies compared to MySQL. Having successfully migrated a rich customer-facing application
suite at the heart of Google's ad business to F1, with no downtime, we will describe how we restructured schema and applications to largely hide this increased latency from external users. The distributed nature of F1 also allows it to scale easily and to support significantly higher throughput for batch workloads than a traditional RDBMS.
With F1, we have built a novel hybrid system that combines the scalability, fault tolerance, transparent sharding, and cost benefits so far available only in ``NoSQL" systems with the usability, familiarity, and transactional guarantees expected from an RDBMS.

Oracle In-Database Hadoop: When MapReduce Meets RDBMS

Xueyuan Su, Yale University; Garret Swart, Oracle

Big data is the tar sands of the data world: vast reserves of raw gritty data whose valuable information content can only be extracted at great cost. MapReduce is a popular parallel programming paradigm well suited to the programmatic extraction and analysis of information from these unstructured Big Data reserves. The Apache Hadoop implementation of MapReduce has become an important player in this market due to its ability to exploit large networks of inexpensive servers. The increasing importance of unstructured data has led to the interest in MapReduce and its Apache Hadoop implementation, which has led to the interest of data processing vendors in supporting this programming style.
Oracle RDBMS has had support for the MapReduce paradigm for many years through the mechanism of user defined pipelined table functions and aggregation objects. However, such support has not been Hadoop source compatible. Native Hadoop programs needed to be rewritten before becoming usable in this framework. The ability to run Hadoop programs inside the Oracle database provides a versatile solution to database users, allowing them use programming skills they may already possess and to exploit the growing Hadoop eco-system.
In this paper, we describe a prototype of Oracle In-Database Hadoop that supports the running of native Hadoop applications written in Java. This implementation executes Hadoop applications using the efficient parallel capabilities of the Oracle database and a subset of the Apache Hadoop infrastructure. This system's target audience includes both SQL and Hadoop users. We discuss the architecture and design, and in particular, demonstrate how MapReduce functionalities are seamlessly integrated within SQL queries. We also share our experience in building such a system within Oracle database and follow-on topics that we think are promising areas for exploration.

SIGMOD Industry 4: Big Data

TAO: How Facebook Serves the Social Graph

Venkateshwaran Venkataramani, Facebook; Zach Amsden, Facebook; Nathan Bronson, Facebook; George Cabrera III, Facebook; Prasad Chakka, Facebook; Peter Dimov, Facebook; Hui Ding, Facebook; Jack Ferris, Facebook; Anthony Giardullo, Facebook; Jeremy Hoon, Facebook; Sachin Kulkarni, Facebook; Nathan Lawrence, Facebook; Mark Marchukov, Facebook; Dmitri Petrov, Facebook; Lovro Puzar, Facebook

Over 800 million people around the world share their social interactions with friends on Facebook, providing a rich body of information referred to as the social graph. In this talk, I describe how we model and serve this graph. Our model uses typed nodes (fbobjects) and edges (associations) to express the relationships and actions that happen on Facebook. We access the graph via a simple API that provides queries over the set of same-typed associations leaving an object. We have found this API to be both sufficiently expressive and amenable to a scalable implementation. In the last segment of the talk I describe the design of TAO, our graph data store. TAO is a distributed implementation of the fbobject and association API that has been serving production traffic at Facebook for more than 2 years.

Large-Scale Machine Learning at Twitter

Jimmy Lin, Twitter; Alex Kolcz, Twitter

The success of data-driven solutions to difficult problems, along with the dropping costs of storing and processing massive amounts of data, has led to growing interest in large-scale machine learning. This paper presents a case study of Twitter's integration of machine learning tools into its existing Hadoop-based, Pig-centric analytics platform. We begin with an overview of this platform, which handles "traditional" data warehousing and business intelligence tasks for the organization. The core of this work lies in recent Pig extensions to provide predictive analytics capabilities that incorporate machine learning, focused specifically on supervised classification. In particular, we have identified stochastic gradient descent techniques for online learning and ensemble methods as being highly amenable to scaling out to large amounts of data. In our deployed solution, common machine learning tasks such as data sampling, feature generation, training, and testing can be accomplished directly in Pig, via carefully crafted loaders, storage functions, and user-defined functions. This means that machine learning is just another Pig script, which allows seamless integration with existing infrastructure for data management, scheduling, and monitoring in a production environment, as well as access to rich libraries of user-defined functions and the materialized output of other scripts.

Recurring Job Optimization in Scope

Nicolás Bruno, Microsoft; Sameer Agarwal, Microsoft; Srikanth Kandula, Microsoft; Bing Shi, Microsoft; Ming-Chuan Wu, Microsoft; Jingren Zhou, Microsoft

SIGMOD Industry 5: Data Integration and Analytics

Dynamic Workload-Driven Data Integration in Tableau

Kristi Morton, University of Washington; Ross Bunker, Tableau Software; Jock Mackinlay, Tableau Software; Robert Morton, Tableau Software; Chris Stolte, Tableau Software

Tableau is a commercial business intelligence (BI) software tool that supports interactive, visual analysis of data. Armed with a visual interface to data and a focus on usability, Tableau enables a wide audience of end-users to gain insight into their datasets. The user experience is a fluid process of interaction in which exploring and visualizing data takes just a few simple drag-and-drop operations (no programming or DB experience necessary). In this context of exploratory, ad-hoc visual analysis, we describe a novel approach to integrating large, heterogeneous data sources. We present a new feature in Tableau called *data blending*, which gives users the ability to create data visualization mashups from structured, heterogeneous data sources dynamically without any upfront integration effort. Users can author visualizations that automatically integrate data from a variety of sources, including data warehouses, data marts, text files, spreadsheets, and data cubes. Because our data blending system is workload driven, we are able to bypass many of the pain-points and uncertainty in creating mediated schemas and schema-mappings in current pay-as-you-go integration systems.

Finding Related Tables

Anish Das Sarma, Google; Lujun Fang, Google; Nitin Gupta, Google; Alon Halevy, Google; Hongrae Lee, Google; Fei Wu, Google; Reynold Xin, Google; Cong Yu, Google

We consider the problem of finding related tables in a large corpus of heterogenous tables. Detecting related tables provides users a powerful tool for enhancing their tables with additional data and enables effective reuse of available public data. Our first contribution is a framework that captures several types of relatedness, including tables that are candidates for joins and

tables that are candidates for union. Our second contribution is a set of algorithms for detecting related tables that can be either unioned or joined. We describe a set of experiments that demonstrate that our algorithms produce highly related tables. We also show that we can often improve the results of table search by pulling up tables that are ranked much lower based on their relatedness to top-ranked tables. Finally, we describe how to scale up our algorithms and show the results of running it on a corpus of over a million tables extracted from Wikipedia.

Optimizing Analytic Data Flows for Multiple Execution Engines

Alkis Simitsis, HP Labs; Kevin Wilkinson, HP Labs; Malu Castellanos, HP Labs; Umeshwar Dayal, HP Labs

Next generation business intelligence involves data flows that span different execution engines, contain complex functionality like data/text analytics, machine learning operations, and need to be optimized against various objectives. Creating correct analytic data flows in such an environment is a challenging task and is both labor-intensive and time-consuming. Optimizing these flows is currently an ad-hoc process where the result is largely dependent on the abilities and experience of the flow designer. Our previous work addressed analytic flow optimization for multiple objectives over a single execution engine. This paper focuses on optimizing flows for a single objective, namely performance, over multiple execution engines. We consider flows that span a DBMS, a Map-Reduce engine, and an orchestration engine (e.g., an ETL tool or scripting language). This configuration is emerging as a common paradigm used to combine analysis of unstructured data with analysis of structured data (e.g., NoSQL plus SQL). We present flow transformations that model data shipping, function shipping, and operation decomposition and we describe how flow graphs are generated for multiple engines. Performance results for various configurations demonstrate the benefit of optimization.

SIGMOD Industry 6: Query Processing and War Stories

CloudRAMSort: Fast and Efficient Large-Scale Distributed RAM Sort on Shared-Nothing Cluster

Changkyu Kim, Intel Labs; Jongsoo Park, Intel Labs; Nadathur Satish, Intel Labs; Hongrae Lee, Google Research; Pradeep Dubey, Intel Labs; Jatin Chhugani, Intel Labs

Sorting is a fundamental kernel used in many database operations. The total memory available across cloud computers is now sufficient to store even hundreds of terabytes of data in-memory. Applications requiring high-speed data analysis typically use in-memory sorting. The two most important factors in designing a high-speed in-memory sorting system are the single-node sorting performance and inter-node communication.

In this paper, we present CloudRAMSort, a fast and efficient system for large-scale distributed sorting on shared-nothing clusters. CloudRAMSort performs multi-node optimizations by carefully overlapping computation with inter-node communication. The system uses a dynamic multi-stage random sampling approach for improved load-balancing between nodes. CloudRAMSort maximizes per-node efficiency by exploiting modern architectural features such as multiple cores and SIMD (Single-Instruction Multiple Data) units. This holistic combination results in the highest performing sorting performance on distributed shared-nothing platforms. CloudRAMSort sorts 1 Terabyte (TB) of data in 4.6 seconds on a 256-node Xeon X5680 cluster called the Intel Endeavor system. CloudRAMSort also performs well on heavily skewed input distributions, sorting 1 TB of data generated using Zipf distribution in less than 5 seconds. We also provide a detailed analytical model that accurately projects (within avg. 7%) the performance of CloudRAMSort with varying tuple sizes and interconnect bandwidths. Our analytical model serves as a useful tool to analyze performance bottlenecks on current systems and project performance with future architectural advances.

With architectural trends of increasing number of cores, bandwidth, SIMD width, cache-sizes, and interconnect bandwidth, we believe CloudRAMSort would be the system of choice for distributed sorting of large-scale in-memory data of current and future systems

Adaptive Optimizations of Recursive Queries in Teradata

Ahmad Ghazal, Teradata; Dawit Seid, Teradata; Alain Crolotte, Teradata; Mohammed Al-Kateb, Teradata

Recursive queries were introduced as part of ANSI SQL 99 to support processing of hierarchical data typical of air flight schedules, bill-of-materials, data cube dimension hierarchies, and ancestor-descendant information (e.g. XML data stored in relations). Recently, recursive queries have also found extensive use in web data analysis such as social network and click stream data. Teradata implemented recursive queries in V2R6 using static plans whereby a query is executed in multiple iterations, each iteration corresponding to one level of the recursion. Such a static planning strategy may not be optimal since the demographics of intermediate results from recursive iterations often vary to a great extent. Gathering feedback at each iteration could address this problem by providing size estimates to the optimizer which, in turn, can produce an execution plan for the next iteration. However, such a full feedback scheme suffers from lack of pipelining and the inability to exploit global optimizations across the different recursion iterations. In this paper, we propose adaptive optimization techniques that avoid the issues with static as well as full feedback optimization approaches. Our approach employs a mix of multi-iteration pre-planning and dynamic feedback techniques which are generally applicable to any recursive query

implementation in an RDBMS. We also validated the effectiveness of our proposed techniques by conducting experiments on a prototype implementation using a real-life social network data from the FriendFeed online blogging service.

From X100 to Vectorwise: Opportunities, Challenges and Things Most Researchers Do Not Think About

Marcin Zukowski, Actian; Peter Boncz, CWI

In 2008 a group of researchers behind the X100 database kernel created Vectorwise: a spin-o which together with the Actian corporation (previously Ingres) worked on bringing this technology to the market. Today, Vectorwise is a popular product and one of the examples of conversion of a research prototype into successful commercial software. We describe here some of the interesting aspects of the work performed by the Vectorwise development team in the process, and discuss the opportunities and challenges resulting from the decision of integrating a prototype-quality kernel with Ingres, an established commercial product. We also discuss how requirements coming from reallife scenarios sometimes clashed with design choices and simplications often found in research projects, and how Vectorwise team addressed some of of them.

SIGMOD DEMONSTRATION ABSTRACTS

SIGMOD Demonstrations A: Information Extraction, Search, Performance, and Clouds

Automatic Web-Scale Information Extraction

Philip Bohannon, Yahoo! Research; Nilesh Dalvi, Yahoo! Research; Yuval Filmus, University of Toronto; Nori Jacoby, Yahoo!; Sathiya Keerthi, Yahoo! Research; Alok Kirpal, Yahoo! Research

In this demonstration, we showcase the technologies that we are building at Yahoo! for Web-scale Information Extraction. Given any new Website, containing semi-structured information about a pre-specified set of schemas, we show how to populate objects in the corresponding schema by automatically extracting information from the Website.

Just-in-Time Information Extraction using Extraction Views

Amr El-Helw, EMC Corp.; Mina Farid, University of Waterloo; Ihab Ilyas, Qatar Computing Research Institute

Many modern applications involve very large amounts of data that comes from unstructured text documents. The ability to answer structured SQL queries over unstructured data allows for more complex analysis and better insights into that data. Querying unstructured data can be accomplished with the help of information extraction techniques. However, the traditional Extract-Transform-Load approach produces data that is out of date. We introduce extraction views, a way to encapsulate IE systems as black boxes. Queries on text documents are evaluated using these extraction views, which enables the optimizer to apply all well-defined optimization techniques. The optimizer selects the best execution plan using a defined cost model.

ColumbuScout: Towards Building Local Search Engines over Large Databases

Cody Hansen, University of Utah; Feifei Li, University of Utah

In many database applications, search is still executed via form based query interfaces, which are then translated into SQL statements to find matching records. Ranking is usually not implemented unless users have explicitly indicated how to rank the matching records, e.g., in the ascending order of year. Often, this approach is neither intuitive nor user friendly (especially with many search fields in a query form). It also requires application developers to design schema-specific query forms and develop specific programs that understand these forms. In this work, we propose to demonstrate the ColumbuScout system that aims at quickly building and deploying a local search engine over one or more large databases. The ColumbuScout system adopts a search-engine-style approach for searches over local databases. It introduces its own indexing structures and storage designs, to improve its overall efficiency and scalability. We will demonstrate that it is simple for application developers to deploy ColumbuScout over any databases, and ColumbuScout is able to support search engine-like types of search over large databases efficiently and effectively.

Sofia Search: A Tool for Automating Related-Work Search

Behzad Golshan, Boston University; Theodoros Lappas, Boston University; Evimaria Terzi, Boston University

When working on a new project, researchers need to devote a significant amount of time and effort to surveying the relevant literature. This is required in order to gain expertise, evaluate the significance of their work and gain useful insights about a particular scientific domain. While necessary, relevant-work search is also a time-consuming and arduous process, requiring the continuous participation of the user. In this work, we introduce Sofia Search, a tool that fully automates the search and retrieval of the literature related to a topic. Given a seed of papers submitted by the user, Sofia Search searches the Web for candidate related papers, evaluates their relevance to the seed and downloads them for the user. The tool also provides modules for the evaluation and ranking of authors and papers, in the context of the retrieved papers. In the demo, we will demonstrate the functionality of our tool, by allowing users to use it via a simple and intuitive interface.

RACE: Real-Time Applications over Cloud-Edge

Badrish Chandramouli, Microsoft Research; Joris Claessens, Microsoft Research; Suman Nath, Microsoft Research; Ivo Santos, Microsoft Research; Wenchao Zhou, University of Pennsylvania

The Cloud-Edge topology - where multiple smart edge devices such as phones are connected to one another via the Cloud - is becoming ubiquitous. We demonstrate RACE, a novel framework and system for specifying and efficiently executing distributed real-time applications in the Cloud-Edge topology. RACE uses LINQ for StreamInsight to succinctly express a diverse suite of useful real-time applications. Further, it exploits the processing power of edge devices and the Cloud to partition and execute such queries in a distributed manner. RACE features a novel cost-based optimizer that efficiently finds the optimal placement, minimizing global communication cost while handling multi-level join queries and asymmetric network links.

Partiqle: An Elastic SQL Engine over Key-Value Stores

Junichi Tatemura, NEC Laboratories America; Oliver Po, NEC Laboratories America; Wang-Pin Hsiung, NEC Laboratories America; Hakan Hacigümüs, NEC Laboratories America

The demo features Partiqle, a SQL engine over key-value stores as a relational alternative for the recent procedural approaches to support OLTP workloads elastically. Based on our *microsharding* framework, it employs a declarative specification, called *transaction classes*, of constraints applied on the transactions in a workload. We demonstrate use of a transaction class in design and analysis of OLTP workloads. We then demonstrate live-scaling of our fully functioning system on a server cluster.

JustMyFriends: Full SQL, Full Transactional Amenities, and Access Privacy

Arthur Meacham, New York University; Dennis Shasha, New York University

A major obstacle to using Cloud services for many enterprises is the fear that the data will be stolen. Bringing the Cloud in-house is an incomplete solution to the problem because that implies that data center personnel as well as myriad repair personnel must be trusted. An ideal security solution would be to share data among precisely the people who should see it ("my friends") and nobody else.

Encryption might seem to be an easy answer. Each friend could download the data, update it perhaps, and return it to a shared untrusted repository. But such a solution permits no concurrency and therefore no real sharing.

JustMyFriends ensures sharing among friends without revealing unencrypted data to anyone outside of a circle of trust. In fact, non-friends (such as system administrators) see only encrypted blobs being added to a persistent store. JustMyFriends allows data sharing and full transactions. It supports the use of all SQL including stored procedures, updates, and arbitrary queries. Additionally, it provides full access privacy, preventing the host from discovering patterns or correlations in the user's data access behavior.

The demonstration will show how friends in an unnamed government agency can coordinate the management of a spy network in a transactional fashion. Demo visitors will be able to play the roles of station chiefs and/or of troublemakers. As station chiefs, they will write their own transactions and queries, logout, login.

As troublemakers, visitors will be able to play the role of a curious observer, kill client processes, and in general try to disrupt the system.

Dynamic Optimization of Generalized SQL Queries with Horizontal Aggregations

Carlos Ordonez, University of Houston; Javier García-García, UNAM; Zhibo Chen, University of Houston

SQL presents limitations to return aggregations as tables with a horizontal layout. A user generally needs to write separate queries and data definition statements to combine transposition with aggregation. With that motivation in mind, we introduce horizontal aggregations, a complementary class of aggregations to traditional (vertical) SQL aggregations. The SQL syntax extension is minimal and it significantly enhances the expressive power and ease of use of SQL. Our proposed SQL extension blurs the boundary between row values and column names. We present a prototype query optimizer that can evaluate arbitrary nested queries combining filtering, joins and both classes of aggregations. Horizontal aggregations have many applications in ad-hoc querying, OLAP cube processing and data mining. We demonstrate query optimization of horizontal aggregations introduces new research challenges.

ConsAD: A Real-Time Consistency Anomalies Detector

Kamal Zellag, McGill University; Bettina Kemme, McGill University

In this demonstration, we present ConsAD, a tool that detects consistency anomalies for arbitrary multi-tier applications that use lower levels of isolation than serializability. As the application is running, ConsAD detects and quantifies anomalies indicating exactly the transactions and data items involved. Furthermore, it classifies the detected anomalies into patterns showing the business methods involved as well as their occurrence frequency. ConsAD can guide designers to either choose an isolation level for which their application shows few anomalies or change their transaction design to avoid the anomalies. Its graphical interface shows detailed information about detected anomalies as they occur and analyzes their patterns as well as their distribution.

Interactive Performance Monitoring of a Composite OLTP and OLAP Workload

Anja Bog, Hasso Plattner Institute, University of Potsdam; Kai Sachs, SAP AG; Hasso Plattner, Hasso Plattner Institute, University of Potsdam

Online transaction processing (OLTP) and online analytical processing (OLAP) are thought of as two separate domains, despite sharing the same business data to operate on. This is the result of performance impairments encountered in the past when running on the same system, the workloads becoming ever more sophisticated, leading to contradictory optimization in database design. Recent developments in hardware and database systems are bringing forth research prototypes supporting mixed OLTP and OLAP workloads, challenging this separation. At the same time new benchmarks are proposed

to assess these mixed workload systems. In the demonstration, we show an interactive performance monitor and benchmark driver developed for the Composite Benchmark for Transaction Processing and Reporting. The performance monitor allows us to directly determine the impact of changing shares within the workload and to interactively assess behavioral characteristics of different database systems under changing mixed workload conditions.

SIGMOD Demonstrations B: Social- or User-Centered

Sindbad: A Location-Based Social Networking System

Mohamed Sarwat, University of Minnesota; Jie Bao, University of Minnesota; Ahmed Eldawy, University Of Minnesota; Justin Levandoski, Microsoft Research; Amr Magdy, University of Minnesota; Mohamed Mokbel, University of Minnesota

This demo presents Sindbad; a location-based social networking system. Sindbad supports three new services beyond traditional social networking services, namely, location-aware news feed, location-aware recommender, and location-aware ranking. These new services not only consider social relevance for its users, but they also consider spatial relevance. Since location-aware social networking systems have to deal with large number of users, large number of messages, and user mobility, efficiency and scalability are important issues. To this end, Sindbad encapsulates its three main services inside the query processing engine of PostgreSQL. Usage and internal functionality of Sindbad, implemented with PostgreSQL and Google Maps API, are demonstrated through user (i.e., web/phone) and system analyzer GUI interfaces, respectively.

MAQSA: A System for Social Analytics on News

Sihem Amer-Yahia, Qatar Computing Research Institute; Samreen Anjum, Qatar Computing Research Institute; Amira Ghenai, Qatar Computing Research Institute; Aysha Siddique, Qatar Computing Research Institute; Sofiane Abbar, Qatar Computing Research Institute; Sam Madden, MIT; Adam Marcus, MIT; Mohammed El-Haddad; Al Jazeera Network

We present MAQSA, a system for social analytics on news. MAQSA provides an interactive topic-centric dashboard that summarizes news articles and social activity (e.g., comments and tweets) around them. MAQSA helps editors and publishers in newsrooms understand user engagement and audience sentiment evolution on various topics of interest. It also helps news consumers explore public reaction on articles relevant to a topic and refine their exploration via related entities, topics, articles and tweets. Given a topic, e.g., ``Gulf Oil Spill," or ``The Arab Spring", MAQSA combines three key dimensions: time, geographic location, and topic to generate a detailed activity dashboard around relevant articles. The dashboard contains an annotated comment timeline and a social graph of comments. It utilizes commenters' locations to build maps of comment sentiment and topics by region of the world. Finally, to facilitate exploration, MAQSA provides listings of related entities, articles, and tweets. It algorithmically processes large collections of articles and tweets, and enables the dynamic specification of topics and dates for exploration. In this demo, participants will be invited to explore the social dynamics around articles on oil spills, the Libyan revolution, and the Arab Spring. In addition, participants will be able to define and explore their own topics dynamically.

Surfacing Time-Critical Insights from Social Media

Bogdan Alexe, IBM Almaden Research Center; Mauricio Hernandez, IBM Almaden Research Center; Kirsten Hildrum, IBM T. J. Watson Research Center; Rajasekar Krishnamurthy, IBM Almaden Research Center; Georgia Koutrika, IBM Almaden Research Center; Meenakshi Nagarajan, IBM Almaden Research Center; Haggai Roitman, IBM Research, Haifa; Michal Shmueli-Scheuer, IBM Research, Haifa; Ioana Stanoi, IBM Almaden Research Center; Chitra Venkatramani, IBM T. J. Watson Research Center; Rohit Wagle, IBM T. J. Watson Research Center

We propose to demonstrate an end-to-end framework for leveraging time-sensitive and critical social media information for businesses. More specifically, we focus on identifying, structuring, integrating, and exposing timely insights that are essential to marketing services and monitoring reputation over social media. Our system includes components for information extraction from text, entity resolution and integration, analytics, and a user interface.

Taagle: Efficient, Personalized Search in Collaborative Tagging Networks

Silviu Maniu, Télécom ParisTech, CNRS LTCI; Bogdan Cautis, Télécom ParisTech, CNRS LTCI

We demonstrate the Taagle system for top-k retrieval in social tagging systems (also known as folksonomies). The general setting is the following: users form a weighted social network, which may reflect friendship, similarity, or trust; items from a public pool of items (e.g., URLs, blogs, photos, documents) are tagged by users with keywords; users search for the top-k items having certain tags. Going beyond a classic search paradigm where data is decoupled from the users querying it, users can now act both as producers and seekers of information. Hence finding the most relevant items in response to a query should be done in a network-aware manner: items tagged by users who are closer (more similar) to the seeker should be given more weight than items tagged by distant users.

We illustrate with Taagle novel algorithms and a general approach that has the potential to scale to current applications, in an online context where the social network, the tagging data and even the seekers' search ingredients can change at any moment. We also illustrate possible design choices for providing users a fully-personalized and customizable search interface. By this interface, they can calibrate how social proximity is computed (for example, with respect to similarity in tagging actions), how much weight the social score of tagging actions should have in the result build-up, or the criteria by which the user network should be explored. In order to further reduce running time, seekers are given the possibility to chose between exact or approximate answers, and can benefit from cached results of previous queries (materialized views).

PrefDB: Bringing Preferences Closer to the DBMS

Anastasios Arvanitis, National Technical University of Athens; Georgia Koutrika, IBM Almaden Research Center

In this demonstration, we present a preference-aware relational query answering system, termed *PrefDB*. The key novelty of *PrefDB* is the use of an extended relational data model and algebra that allow expressing different flavors of preferential queries. Furthermore, unlike existing approaches that either treat the DBMS as a black box or require modifications of the database core, *PrefDB*'s hybrid implementation enables operator-level query optimizations without being obtrusive to the database engine. We showcase the flexibility and efficiency of *PrefDB* using *PrefDBAdmin*, a graphical tool that we have built aiming at assisting application designers in the task of building, testing and tuning queries with preferences.

Auto-Completion Learning for XML

Serge Abiteboul, Collège de France, INRIA Saclay, ENS Cachan; Yael Amsterdamer, Tel Aviv University; Tova Milo, Tel Aviv University; Pierre Senellart, Télécom ParisTech, CNRS LTCI

Editing an XML document manually is a complicated task. While many XML editors exist in the market, we argue that some important functionalities are missing in all of them. Our goal is to makes the editing task simpler and faster. We present ALEX (Auto-completion Learning Editor for XML), an editor that assists the users by providing intelligent auto-completion suggestions. These suggestions are adapted to the user needs, simply by feeding ALEX with a set of example XML documents to learn from. The suggestions are also guaranteed to be compliant with a given XML schema, possibly including integrity constraints. To fulfill this challenging goal, we rely on novel, theoretical foundations by us and others, which are combined here in a system for the first time.

Logos: A System for Translating Queries into Narratives

Andreas Kokkalis, University of Athens; Panagiotis Vagenas, University of Athens; Alexandros Zervakis, University of Athens; Alkis Simitsis, HP Labs; Georgia Koutrika, IBM Almaden Research Center; Yannis Ioannidis, University of Athens

This paper presents Logos, a system that provides natural lan-guage translations for relational queries expressed in SQL. Our translation mechanism is based on a graph-based approach to the query translation problem. We represent various forms of structured queries as directed graphs and we annotate the graph edges with template labels using an extensible template mechanism. Logos uses different graph traversal strategies for efficiently exploring these graphs and composing textual query descriptions. The audience may interactively explore Logos using various database schemata and issuing either sample or ad hoc queries.

PAnG-Finding Patterns in Annotation Graphs

Philip Anderson, University of Maryland; Andreas Thor, University of Maryland; Joseph Benik, University of Maryland; Louiqa Raschid, University of Maryland; María Esther Vidal, Universidad Simón Bolívar

Annotation graph datasets are a natural representation of scientific knowledge. They are common in the life sciences and health sciences, where concepts such as genes, proteins or clinical trials are annotated with controlled vocabulary terms from ontologies. We present a tool, PAnG (Patterns in Annotation Graphs), that is based on a complementary methodology of graph summarization and dense subgraphs. The elements of a graph summary correspond to a pattern and its visualization can provide an explanation of the underlying knowledge. Scientists can use PAnG to develop hypotheses and for exploration.

VizDeck: Self-Organizing Dashboards for Visual Analytics

Alicia Key, University of Washington; Bill Howe, University of Washington; Daniel Perry, University of Washington; Cecilia Aragon, University of Washington

We present VizDeck, a web-based tool for exploratory visual analytics of unorganized relational data. Motivated by collaborations with domain scientists who search for complex patterns in hundreds of data sources simultaneously, VizDeck automatically recommends appropriate visualizations based on the statistical properties of the data and adopts a card game metaphor to help organize the recommended visualizations into interactive visual dashboard applications in seconds with

zero programming. The demonstration allows users to derive, share, and permanently store their own dashboard from hundreds of real science datasets using a production system deployed at the University of Washington.

Kaizen: A Semi-Automatic Index Advisor

Ivo Jimenez, University of California, Santa Cruz; Huascar Sanchez, University of California, Santa Cruz; Quoc Trung Tran, University of California, Santa Cruz; Neoklis Polyzotis, University of California, Santa Cruz

Index tuning; i.e., selecting indexes that are appropriate for the workload to obtain good system performance, is a crucial task for database administrators.
Administrators rely on automated index advisors for this task, but existing advisors work either offline, requiring a-priori knowledge of the workload, or online, taking the administrator out of the picture and assuming total control of the index tuning task. Semi-automatic index tuning is a new paradigm that achieves a middle ground: the advisor analyzes the workload online and provides recommendations tailored to the current workload, and the administrator is able to provide feedback to refine future recommendations. In this demonstration we present Kaizen, an index tuning tool that implements semi-automatic tuning.

SIGMOD Demonstrations C: Analytics

Shark: Fast Data Analysis Using Coarse-Grained Distributed Memory

Cliff Engle, University of California, Berkeley; Antonio Lupher, University of California, Berkeley; Reynold Xin, University of California, Berkeley; Matei Zaharia, University of California, Berkeley; Michael Franklin, University of California, Berkeley; Scott Shenker, University of California, Berkeley; Ion Stoica, University of California, Berkeley

Shark is a research data analysis system built on a novel coarse-grained distributed shared-memory abstraction. Shark marries query processing with deep data analysis, providing a unified system for easy data manipulation using SQL and pushing sophisticated analysis closer to data. It scales to thousands of nodes in a fault-tolerant manner. Shark can answer queries 40X faster than Apache Hive and run machine learning programs 25X faster than MapReduce programs in Apache Hadoop on large datasets.

Exploiting MapReduce-Based Similarity Joins

Yasin N. Silva, Arizona State University; Jason M. Reed, Arizona State University

Cloud enabled systems have become a crucial component to efficiently process and analyze massive amounts of data. One of the key data processing and analysis operations is the Similarity Join, which retrieves all data pairs whose distances are smaller than a pre-defined threshold μ. Even though multiple algorithms and implementation techniques have been proposed for Similarity Joins, very little work has addressed the study of Similarity Joins for cloud systems. This paper presents MRSimJoin, a multi-round MapReduce based algorithm to efficiently solve the Similarity Join problem. MRSimJoin efficiently partitions and distributes the data until the subsets are small enough to be processed in a single node. The proposed algorithm is general enough to be used with data that lies in any metric space. We have implemented MRSimJoin in Hadoop, a highly used open-source cloud system. We show how this operation can be used in multiple real-world data analysis scenarios with multiple data types and distance functions. Particularly, we show the use of MRSimJoin to identify similar images represented as feature vectors, and similar publications in a bibliographic database. We also show how MRSimJoin scales in each scenario when important parameters, e.g., μ, data size and number of cluster nodes, increase. We demonstrate the execution of MRSimJoin queries using an Amazon Elastic Compute Cloud (EC2) cluster.

GLADE: Big Data Analytics Made Easy

Yu Cheng, University of California, Merced; Chengjie Qin, University of California, Merced; Florin Rusu, University of California, Merced

We present GLADE, a scalable distributed system for large scale data analytics. GLADE takes analytical functions expressed through the User-Defined Aggregate (UDA) interface and executes them efficiently on the input data. The entire computation is encapsulated in a single class which requires the definition of four methods. The runtime takes the user code and executes it right near the data by taking full advantage of the parallelism available inside a single machine as well as across a cluster of computing nodes.
The demonstration has two goals. First, it presents the architecture of GLADE and how processing is done by using a series of analytical functions. Second, it compares GLADE with two different classes of systems for data analytics: a relational database (PostgreSQL) enhanced with UDAs and Map-Reduce (Hadoop). We show how the analytical functions are coded into each of these systems (for Map-Reduce, we use both Java code as well as Pig Latin) and compare their expressiveness, scalability, and running time efficiency.

ReStore: Reusing Results of MapReduce Jobs in Pig

Iman Elghandour, University of Waterloo; Ashraf Aboulnaga, University of Waterloo

Analyzing large scale data has become an important activity for many organizations, and is now facilitated by the MapReduce programming and execution model and its implementations, most notably Hadoop. Query languages such as Pig Latin, Hive, and Jaql make it simpler for users to express complex analysis tasks, and the compilers of these languages translate these complex tasks into workflows of MapReduce jobs. Each job in these workflows reads its input from the distributed file system used by the MapReduce system (e.g., HDFS in the case of Hadoop) and produces output that is stored in this distributed file system. This output is then read as input by the next job in the workflow. The current practice is to delete these intermediate results from the distributed file system at the end of executing the workflow. It would be more useful if these intermediate results can be stored and reused in future workflows. We demonstrate ReStore, an extension to Pig that enables it to manage storage and reuse of intermediate results of the MapReduce workflows executed in the Pig data analysis system. ReStore matches input workflows of MapReduce jobs with previously executed jobs and rewrites these workflows to reuse the stored results of the matched jobs. ReStore also creates additional reuse opportunities by materializing and reserving the output of query execution operators that are executed within a MapReduce job. In this demonstration we showcase the MapReduce jobs and sub-jobs recommended by ReStore for a given Pig query, the rewriting of input queries to reuse stored intermediate results, and a what-if analysis of the effectiveness of reusing stored outputs of previously executed jobs.

Clydesdale: Structured Data Processing on Hadoop

Andrey Balmin, IBM Almaden Research Center; Tim Kaldewey, IBM Almaden Research Center; Sandeep Tata, IBM Almaden Research Center

There have been several recent proposals modifying Hadoop, radically changing the storage organization or query processing techniques to obtain good performance for structured data processing. We will showcase Clydesdale, a research prototype for structured data processing on Hadoop that can achieve dramatic performance improvements over existing solutions, without any changes to the underlying MapReduce implementation. Clydesdale achieves this through a novel synthesis of several techniques from the database literature and carefully adapting them to the Hadoop environment. On the star schema benchmark, we show that Clydesdale is on average 38x faster than Hive, the dominant approach for structured data processing on Hadoop today. To the best of our knowledge, Clydesdale is the fastest solution for processing workloads on structured data sets that fit a star schema on Hadoop. Attendees will be able to run queries on the data from the star schema benchmark on a remote Hadoop cluster with Clydesdale and Hive installed, and get a breakdown of the time taken to execute the query. Attendees will also be able to pose their own queries using ClyQL -- a novel embedded DSL in Scala that can be used to rapidly prototype star join queries. With this demonstration, we hope to convince the attendees that unlike previously thought, Hadoop can indeed efficiently support structured data processing.

Tiresias: A Demonstration of How-To Queries

Alexandra Meliou, University of Washington; Yisong Song, University of Washington; Dan Suciu, University of Washington

In this demo, we will present Tiresias, the first how-to query engine. How-to queries represent fundamental data analysis questions of the form: ``How should the input change in order to achieve the desired output''. They exemplify an important Reverse Data Management problem: solving constrained optimization problems over data residing in a DBMS.
Tiresias, named after the mythical oracle of Thebes, has complex under-workings, but includes a simple interface that allows users to load datasets and interactively design optimization problems by simply selecting actions, key performance indicators, and objectives. The user choices are translated into a declarative query, which is then processed by Tiresias and translated into a Mixed Integer Program: we then use an MIP solver to find a solution. The solution is then presented to the user as an interactive data instance. The user can provide feedback by rejecting certain tuples and/or values. Then, based on the user feedback, Tiresias automatically refines the how-to query and presents a new set of results.

AstroShelf: Understanding the Universe Through Scalable Navigation of a Galaxy of Annotations

Panayiotis Neophytou, University of Pittsburgh; Roxana Gheorghiu, University of Pittsburgh; Rebecca Hachey, University of Pittsburgh; Timothy Luciani, University of Pittsburgh; Di Bao, University of Pittsburgh; Alexandros Labrinidis, University of Pittsburgh; Elisabeta G. Marai, University of Pittsburgh; Panos K. Chrysanthis, University of Pittsburgh

This demo presents AstroShelf, our on-going effort to enable astrophysicists to collaboratively investigate celestial objects using data originating from multiple sky surveys, hosted at different sites. The AstroShelf platform combines database and data stream, workflow and visualization technologies to provide a means for querying and displaying telescope images (in a Google Sky manner), visualizations of spectrum data, and for managing annotations. In addition to the user interface, AstroShelf supports a programmatic interface (available as a web service), which allows astrophysicists to incorporate functionality from AstroShelf in their own programs. A key feature is Live Annotations which is the detection and delivery

of events or annotations to users in real-time, based on their profiles. We demonstrate the capabilities of AstroShelf through real end-user exploration scenarios (with participation from "stargazers" in the audience), in the presence of simulated annotation workloads executed through web services.

OPAvion: Mining and Visualization in Large Graphs

Leman Akoglu, Carnegie Mellon University; Duen Horng Chau, Carnegie Mellon University; U Kang, Carnegie Mellon University; Danai Koutra, Carnegie Mellon University; Christos Faloutsos, Carnegie Mellon University

Given a large graph with millions or billions of nodes and edges, like a who-follows-whom Twitter graph, how do we scalably compute its statistics, summarize its patterns, spot anomalies, visualize and make sense of it? We present OPAvion, a graph mining system that provides a scalable, interactive workflow to accomplish these analysis tasks.
OPAvion consists of three modules:

(1) The *Summarization* module (Pegasus) operates off-line on massive, disk-resident graphs and computes graph statistics, like PageRank scores, connected components, degree distribution, triangles, etc.;

(2) The *Anomaly Detection* module (OddBall) uses graph statistics to mine patterns and spot anomalies, such as nodes with many contacts but few interactions with them (possibly telemarketers);
(3) The *Interactive Visualization* module (Apolo) lets users incrementally explore the graph, starting with their chosen nodes or the flagged anomalous nodes; then users can expand to the nodes' vicinities, label them into categories, and thus interactively navigate the interesting parts of the graph. In our demonstration, we invite our audience to interact with OPAvion and try out its core capabilities on the Stack Overflow Q&A graph that describes over 6 million questions and answers among 650K users.

CloudAlloc: A Monitoring and Reservation System for Compute Clusters

Enrico Iori, University of Trento; Alkis Simitsis, HP Labs; Themis Palpanas, University of Trento; Kevin Wilkinson, HP Labs; Stavros Harizopoulos, Nou Data

Cloud computing has emerged as a promising environment capable of providing flexibility, scalability, elasticity, fail-over mechanisms, high availability, and other important features to applications. Compute clusters are relatively easy to create and use, but tools to effectively share cluster resources are lacking. CloudAlloc addresses this problem and schedules workloads to cluster resources using allocation algorithms that can be easily changed according to the objectives of the enterprise. It also monitors resource utilization and thus, provides accountability for actual usage. CloudAlloc is a lightweight, flexible, easy-to-use tool for cluster resource allocation that has also proved useful as a research platform. We demonstrate its features, and also discuss its allocation algorithms that minimize power usage. CloudAlloc was implemented and is in use at HP Labs.

TIRAMOLA: Elastic NoSQL Provisioning through a Cloud Management Platform

Ioannis Konstantinou, National Technical University of Athens; Evangelos Angelou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Christina Boumpouka, National Technical University of Athens; Nectarios Koziris, National Technical University of Athens; Spyros Sioutas, Ionian University

NoSQL databases focus on analytical processing of large scale datasets, offering increased scalability over commodity hardware. One of their strongest features is elasticity, which allows for fairly portioned premiums and high-quality performance. Yet, the process of adaptive expansion and contraction of resources usually involves a lot of manual effort, often requiring the definition of the conditions for scaling up or down to be provided by the users. To date, there exists no open-source system for automatic resizing of NoSQL clusters.
In this demonstration, we present *TIRAMOLA*, a modular, cloud-enabled framework for monitoring and adaptively resizing NoSQL clusters. Our system incorporates a decision-making module which allows for optimal cluster resize actions in order to maximize any quantifiable reward function provided together with life-long adaptation to workload or infrastructural changes.
The audience will be able to initiate HBase clusters of various sizes and apply varying workloads through multiple YCSB clients. The attendees will be able to watch, in real-time, the system perform automatic VM additions and removals as well as how cluster performance metrics change relative to the optimization parameters of their choice.

UNDERGRADUATE POSTER COMPETITION

Declarative Web Application Development: Encapsulating Dynamic JavaScript Widgets

Robert Bolton, University of California, San Diego; David Ing, University of California, San Diego; Christopher Rebert, University of California, San Diego; Kristina Lam Thai, University of California, San Diego

The development of modern, highly interactive AJAX Web applications that enable dynamic visualization of data requires writing a great deal of tedious plumbing code to interface data between browser-based DOM and AJAX components, the application server, and the SQL database. Worse, each of these layers utilizes a different language. Further, much code is needed to keep the page and application states in sync using an imperative paradigm, which hurts simplicity. These factors result in a frustrating experience for today's Web developer. The FORWARD Project aims to alleviate this frustration by enabling pages that are rendered views, in the SQL sense of view.

Our work in the project has led to a highly declarative approach whereby JavaScript/AJAX UI widgets automatically render views over the application state (database + session data + page data) without requiring the developer to tediously code how changes to the application state lead to invocation of the componentsupdate methods.
In contrast to conventional Web application development approaches, a FORWARD application involves only two languages, both declarative: an extended version of SQL, and an XML-based language for configuration and orchestration. The framework automatically handles efficient exchange of user input and changes to the underlying data, and updates the application state accordingly. The developer does not need to write any JavaScript or explicit updating code themselves. On the client side, FORWARD units wrap widgets using JavaScript to collect user input, directly display data, and reflect server-side updates to the data. On the server side, units contain Java code necessary to expose their functionality to the FORWARD framework and define their XML configuration representation.
Our demo consists of a dynamically rendered webpage which internally uses AJAX to update a Google Maps widget that shows location markers for current Groupon deals in a specified area. It will illustrate that our SQL-driven approach makes this kind of rich dynamic webpage easy to write, with significant improvements in simplicity, brevity, and development time, while still providing the quality experience expected from top AJAX components. The amount of plumbing code is significantly reduced, enhancing the experience of AJAX Web application developers.

Towards Scalable Summarization and Visualization of Large Text Corpora

Tyler Sliwkanich, University of Alberta; Douglas Schneider, University of Alberta; Aaron Yong, University of Alberta; Mitchell Home, University of Alberta; Denilson Barbosa, University of Alberta

Society is awash with problems requiring the analysis of vast quantities of text and data. From detecting flu trends out of twitter conversations to finding scholarly works answering specific questions, we rely more and more on computers to process text for us. Text analytics is the application of computational, mathematical, and statistical models to derive information from large quantities of data coming primarily as text. Our project provides fast and effective text-analytics tools for large document collections, such as the blogosphere. We use natural language processing and database techniques to extract, collect, analyze, visualize, and archive information extracted from text. We focus on discovering relationships between entities (people, places, organizations, etc.) mentioned in one or more sources (blog posts or news articles). We built a custom solution using mostly off-the-shelf, open-source tools to provide a scalable platform for users to search and analyze large text corpora. Currently, we provide two main outlets for users to discover these relations: (1) full-text search over the documents and (2) graph visualizations of the entities and their relationships. This provides the user with succinct and easily digestible information gleaned from the corpus as a whole. For example, we can easily pose queries like which companies were bought by Google? as entity:google relation:bought. The extracted data is stored on a combination of the noSQL database CouchDB and Apache's Lucene. This combination is justified as our work-flow consists of offline batch insertions with almost no updates. Because we support specialized queries, we can forgo the flexibility of traditional SQL solutions and materialize all necessary indices, which are used to quickly query large amounts of de-normalized data using MapReduce. Lucene provides a flexible and powerful query syntax to yield relevant ranked results to the user. Moreover, its indices are synchronized by a process subscribed to the list of database changes published by CouchDB. The graph visualizations rely on CouchDB's ability to export the data in any format: we currently use a customized graph visualization relying on XML data. Finally, we use memcached to further improve the performance, especially for queries involving popular entities.

Reducing Cache Misses in Hash Join Probing Phase By Pre-Sorting Strategy

Gi-Hwan Oh, SungKyunKwan University; Jae-Myung Kim, SungKyunKwan University; Woon-Hak Kang, SungKyunKwan University; Sang-Won Lee, SungKyunKwan University

Recently, several studies on multi-core cache-aware hash join have been carried out. In particular, the work of Blanas has shown that rather simple no-partitioning hash join can outperform the work of Kim. Meanwhile, the simple but best performing hash join of Blanas still experiences severe cache misses in probing phase. Because the key values of tuples in outer relation are not sorted or clustered, each outer record has different hashed key value and thus accesses the different hash bucket. Since the size of hash table of inner table is usually much larger than that of the CPU cache, it is highly probable that the reference to hash bucket of inner table by each outer record would encounter cache miss. To reduce the cache misses in hash join probing phase, we propose a new join algorithm, Sorted Probing (in short, SP), which pre-sorts the hashed key values of outer table of hash join so that the access to the hash bucket of inner table has strong temporal locality, thus minimizing the cache misses during the probing phase. As an optimization technique of sorting, we used the cache-aware AlphaSort technique, which extracts the key from each record of data set to be sorted and its pointer, and then sorts the pairs of (key, rec_ptr). For performance evaluation, we used two hash join algorithms from Blanas' work, no partitioning(NP) and independent partitioning(IP) in a standard C++ program, provided by Blanas. Also, we implemented the AlphaSort and added it before each probing phase of NP and IP, and we call each algorithm as NP+SP and IP+SP. For syntactic workload, IP+SP outperforms all other algorithms: IP+SP is faster than other altorithms up to 30%.

DP-tree: Indexing Multi-Dimensional Data under Differential Privacy

Shangfu Peng, Shanghai Jiao Tong University; Yin Yang, Advanced Digital Sciences Center; Zhenjie Zhang, Advanced Digital Sciences Center; Marianne Winslett, Advanced Digital Sciences Center; Yong Yu, Shanghai Jiao Tong University

e-differential privacy (e-DP) is a strong and rigorous scheme for protecting individualsprivacy while releasing useful statistical information. The main idea is to inject random noise into the results of statistical queries, such that the existence of any single record has negligible impact on the distributions of query results. The accuracy of such randomized results depends heavily upon the query processing technique, which has been an active research topic in recent years. So far, most existing methods focus on 1-dimensional queries. The only work that handles multi-dimensional query processing under e-DP is [1], which indexes the sensitive data using variants of the quad-tree and the k-d-tree. As we point out in this paper, these structures are inherently suboptimal for answering queries under e-DP. Consequently, the solutions in [1] suffer from several serious drawbacks, including limited and unstable query accuracy, as well as bias towards certain types of queries.

Motivated by this, we propose the DP-tree, a novel index structure for multi-dimensional query processing under e-DP that eliminates the problems encountered by the methods in [1]. Further, we show that the effectiveness of the DP-tree can be improved using statistical information about the query workload. Extensive experiments using real and synthetic datasets confirm that the DP-tree achieves significantly higher query accuracy than existing methods. Interestingly, an adaptation of the DP-tree also outperforms previous 1D solutions in their restricted scope, by large margins.

Temporal Provenance Discovery in Micro-Blog Message Streams

Zijun Xue, Peking University; Junjie Yao, Peking University; Bin Cui, Peking University

Recent years have witnessed the flourishing increases of micro-blog message applications. Prominent examples include Twitter, Facebook's status, and Sina Weibo in China. Messages in these applications are short (140 characters in a message) and easy to create. The subscription and re-sharing features also make it fairly intuitive to propagate. Micro-blog applications provide abundant information to present world scale user interests and social pulse in an unexpected way. But the precious corpus also brings out the noise and fast changing fragments to prohibit effective understanding and management.

In this work, we propose a micro-blog provenance model to capture temporal connections within micro-blog messages. Here, provenance refers to data origin identification and transformation logging, demonstrating of great value in recent database and workflow systems. The provenance model is used to represent the message development trail and changes explicitly. We select various types of connections in micro-blog applications to identify the provenance. To cope with the real time micro-message deluge, we discuss a novel message grouping approach to encode and maintain the provenance information. A summary index structure is utilized to enable efficient provenance updating. We collect in-coming messages and compare them with an in-memory index to associate them with related ones. The closely related messages form some virtual provenance representation in a coarse granularity. We periodically dump memory values onto disks.

In the actual implementation, we also introduce several adaptive pruning strategies to extend the potential of provenance discovery efficiency. We use the temporal decaying and granularity levels to filter out low chance messages. In the demonstration, we reveal the usefulness of provenance information for rich query retrieval and dynamic message tracking for effective message organization. The real-time collection approach shows advantages over some baselines. Experiments conducted on a real dataset verify the effectiveness and efficiency of our provenance approach. Results show that the

partial-indexing strategy and other restriction ones can maintenance the accuracy at 90% and returning rate at 60% with a reasonable low memory usage. This is the first work towards provenance-based indexing support for micro-blog platforms.

SigSpot: Mining Significant Anomalous Regions from Time-Evolving Networks

Misael Mongiovì, University of California, Santa Barbara; Petko Bogdanov, University of California, Santa Barbara; Razvan Ranca, University of California, Santa Barbara; Ambuj K. Singh, University of California, Santa Barbara; Evangelos E. Papalexakis, Carnegie Mellon University; Christos Faloutsos, Carnegie Mellon University

Anomaly detection in dynamic networks has a rich gamut of application domains, such as road networks, communication networks and water distribution networks. An anomalous event, such as a traffic accident, denial of service attack or a chemical spill, can cause a local shift from normal behavior in the network state that persists over an interval of time. Detecting such anomalous regions of network and time extent in large real-world networks is a challenging task. Existing anomaly detection techniques focus on either the time series associated with individual network edges or on global anomalies that affect the entire network. In order to detect anomalous regions, one needs to consider both the time and the affected network substructure jointly, which brings forth computational challenges due to the combinatorial nature of possible solutions.

We propose the problem of mining all Significant Anomalous Regions (SAR) in time-evolving networks that asks for the discovery of connected temporal subgraphs comprised of edges that significantly deviate from normal in a persistent manner. We propose an optimal Baseline algorithm for the problem and an efficient approximation, called SIGSPOT . Compared to Baseline, SIGSPOT is up to one order of magnitude faster in real data, while achieving less than 10% average relative error rate. In synthetic datasets it is more than 30 times faster than Baseline with 94% accuracy and solves efficiently large instances that are infeasible (more than 10 hours running time) for Baseline. We demonstrate the utility of SIGSPOT for inferring accidents on road networks and study its scalability when detecting anomalies in social, transportation and synthetic evolving networks, spanning up to 1GB.

VRRC: Web Based Tool for Visualization and Recommendation on Co-Authorship Network

Eduardo M. Barbosa, UFMG; Mirella M. Moro, UFMG; Giseli Rabello Lopes, UFRGS; J. Palazzo M. de Oliveira, UFRGS

Scientific studies are usually developed by contributions from different researchers. Analyzing such collaborations is often necessary, for example, when evaluating the quality of a research group. Also, identifying new partnership possibilities within a set of researchers is frequently desired, for example, when looking for partners in foreign countries. Both analysis and identification are not easy tasks, and are usually done manually. This work presents VRRC, a new approach for visualizing recommendations of people within a co-authorship network (i.e., a graph in which nodes represent researchers and edges represent their co-authorships). VRRC input is a publication list from which it extracts the co-authorships. VRRC then recommends which relations could be created or intensified based on metrics designed for evaluating co-authorship networks. Finally, VRRC provides brand new ways to visualize not only the final recommendations but also the intermediate interactions within the network, including: a complete representation of the co-authorship network; an overview of the collaborations evolution over time; and the recommendations for each researcher to initiate or intensify cooperation. Some visualizations are interactive, allowing to filter data by time frame and highlighting specific collaborations. The contributions of our work, compared to the state-of-art, can be summarized as follows: (i) VRRC can be applied to any co-authorship network, it provides both net and recommendation visualizations, it is a Web-based tool and it allows easy sharing of the created visualizations (existing tools do not offer all these features together); (ii) VRRC establishes graphical representations to ease the visualization of its results (traditional approaches present the recommendation results through simple lists or charts); and (iii) with VRRC, the user can identify not only new possible collaborations but also existing cooperation that can be intensified (current recommendation approaches only indicate new collaborations). This work was partially supported by CNPq, Brazil.

Fast Sampling Word Correlations of High Dimensional Text Data

Frank Rosner, Martin-Luther-University Halle-Wittenberg; Alexander Hinneburg, Martin-Luther-University Halle-Wittenberg; Martin Gleditzsch, Unister GmbH; Mathias Priebe, Unister GmbH; Andreas Both, Unister GmbH

Finding correlated words in large document collections is an important ingredient for text analytics. The naïve approach computes the correlations of each word against all other words and filters for highly correlated word pairs. Clearly, this quadratic method cannot be applied to real world scenarios with millions of documents and words. Our main contribution is to transform the task of finding highly correlated word pairs into a word clustering problem that is efficiently solved by locality sensitive hashing (LSH). A key insight of our new method is to note that the empirical Pearson correlation between two words is the cosine of the angle between the centered versions of their word vectors. The angle can be approximated by an LSH scheme. Although centered word vectors are not sparse, the computation of the LSH hash functions can exploit the inherent sparsity of the word data. This leads to an efficient way to detect collisions between centered word vectors having a

small angle and therefore provides a fast algorithm to sample highly correlated word pairs. Our new method based on LSH improves run time complexity of the enhanced naïve algorithm. This algorithm reduces the dimensionality of the word vectors using random projection and approximates correlations by computing cosine similarity on the reduced and centered word vectors. However, this method still has quadratic run time. Our new method replaces the filtering for high correlations in the naïve algorithm with finding hash collisions, which can be done by sorting the hash values of the word vectors. We evaluate the scalability of our new algorithm to large text collections.

PROGRAMMING CONTEST FINALISTS

Team 1
Jung-Sang Ahn (KAIST, South Korea)

Team 2
Amin Allam (KAUST, Saudi Arabia)

Team 3
Guoda Chen (Carnegie Mellon University, USA); Han Liu (KAUST, Saudi Arabia)

Team 4
Ioana Ileana (Telecom ParisTech, France); Karsten Schock (Telecom ParisTech, France)

Team 5
Fuad Jamour (KAUST, Saudi Arabia); Lie Yan (KAUST, Saudi Arabia); Islam al-Masri (KAUST, Saudi Arabia)

SIGMOD TRAVEL AWARDS

Jie Bao	University of Minnesota
Yingyi Bu	University of California, Irvine
Inci Cetindil	University of California, Irvine
Qian Chen	Hong Kong Baptist University
Krzysztof Choromanski	Columbia University
Eli Cortez	Federal University of Amazonas
BOLIN DING	UIUC
Boxiang Dong	Stevens Institute of Tech
Liyue Fan	Emory University
Lujun Fang	University of Michigan
Avrilia Floratou	UW-Madison
Junhao Gan	Sun Yat-Sen University
Roxana Gheorghiu	University of Pittsburgh
Nikos Giatrakos	University of Piraeus
Behzad Golshan	Boston University
Raman Grover	UCI
Benoit Groz	INRIA
Anja Gruenheid	Technical University of Munich
Cody Hansen	University of Utah
Yeye He	University of Wisconsin-Madison
Zengfeng Huang	HKUST
George Konstantinidis	U. of Southern California
Jeffrey LeFevre	UC Santa Cruz
Andres Letelier	PUC
Reut Levi	Tel Aviv University
Zheng Li	UMass Lowell
Erietta Liarou	CWI
Amr Magdy	University of Minnesota
Prashanth Mohan	UC Berkeley
Jayanta Mondal	University of Maryland
Abhishek Mukherji	WPI

Panayiotis Neophytou	University of Pittsburgh
Chongling Nie	ETH Zurich
Giuseppe Ottaviano	Università di Pisa
Enela Pema	UCSC
Thao Pham	University of Pittsburgh
Li Qian	University of Michigan
Karthik Ramachandra	IIT Bombay
Padmashree Ravindra	North Carolina State Universit
Theodoros Rekatsinas	University of Maryland
Miguel Romero	University of Chile
Manish Singh	University of Michigan
Xueyuan Su	Yale University
Abhradeep Guha Thakurta	Pennsylvania State University
Matthew Tucker	University of Nebraska Omaha
Panagiotis Vagenas	University of Athens
Jiannan Wang	Tsinghua University
Jingwen Wang	UMass Lowell
Di Wang	Worcester Polytechnic Institute
Zhewei Wei	HKUST
Wentao Wu	Univ. of Wisconsin-Madison
Mohamed Yakout	Purdue University
Shengqi Yang	UCSB
Yuli Ye	University of Toronto
Wuzhou Zhang	Duke University
Bo Zong	UCSB

PH.D. WORKSHOP PARTICIPANTS

Mohammad Sadoghi (University of Toronto, Canada)

Gregor Endler (University of Erlangen-Nuremberg, Germany)

Ablimit Aji (Emory University, USA)

Samujjwal Bhandari (Texas Tech University, USA)

Pengcheng Xiong (Georgia Institute of Technology, USA)

Katja Losemann (Universität Bayreuth, Germany)

Sebastian Skritek (Vienna University of Technology, Austria)

Eleni Petraki (CWI, Netherlands)

Pei Li (University of Milan – Bicocca, Italy)

Mengmeng Liu (University of Pennsylvania, USA)

Mohamed Sarwat (University of Minnesota, USA)

Filipe Mesquita (University of Alberta, Canada)

CO-LOCATED WORKSHOPS

SIGMOD/PODS Ph.D. Symposium
http://db.uwaterloo.ca/PhD2012/
May 20, 2012
Location: Dunes A-B

The symposium will bring together Ph.D. students working on topics related to the SIGMOD/PODS conference. The workshop offers Ph.D. students the opportunity to present, discuss, and receive feedback on their research in a constructive and international atmosphere. The workshop is be accompanied by prominent professors, researchers and practitioners in the fields of database technology. These accompanying professors participate actively and contribute to the discussions.

SIGMOD 2012 DataBase MEntoring Workshop (DB Me 2012)
http://www.cs.ubc.ca/~rap/dbme2012/
May 20, 2012
Location: Arizona Ballroom VIII

The workshop aims to increase the diversity of the database community through mentoring of students, especially women and underrepresented minorities. This is the second DB Me workshop; the first was held in conjunction with SIGMOD 2010.

2nd ACM SIGMOD Workshop on Databases and Social Networks (DBSocial 2012)
https://sites.google.com/site/dbsocial12/
May 20, 2012
Location: Arizona Ballroom II

The Second ACM SIGMOD Workshop on Databases and Social Networks (DBSocial 2012) is a venue for database research applied to the problems of extraction, querying, and analysis of social networks. DBSocial aims at disseminating results founded on database research and practice that advance the state-of-the-art in the observation, management, and analysis of inherently networked data originating primarily from social phenomena.

DBSocial spans theoretical as well as practical research whose approaches are within the scope of databases and very closely related areas (e.g., data mining and information retrieval), and whose validation are on par with the high standards in the database community.

9th International Workshop on Information Integration on the Web (IIWeb 2012)
http://research.ihost.com/iiweb12/index.html
May 20, 2012
Location: Arizona Ballroom VII

This, the ninth workshop in the IIWeb series, is focused on identifying challenges to be overcome for effectively extracting and integrating knowledge from the Web, enterprise data, and social media. The purpose of this workshop is to bring together researchers working in a variety of areas that are all related to the larger problem of integrating information. This includes research on information extraction, data integration, semantic web, web services, data mining, query execution, and other related areas.

The workshop's focus is on Social and web data integration.

3rd International Workshop on Keyword Search on Structured Data (KEYS 2012)
http://datasearch.ruc.edu.cn/keys2012/
May 20, 2012
Location: Arizona Ballroom VI

Keyword-based web search engines are widely used for searching documents, images, and video. There are also vast collections of structured and semi-structured data on the Web and in enterprises. Traditionally, to access these resources, a user must first learn query languages, schemas, etc. Keyword-based search of such data offers users easier access to the data. The workshop covers also keyword search on spatial data research and micro-blogging content. Geographic objects with associated descriptive texts call for spatial keyword queries that take into account both the locations and textual descriptions of the objects. Micro-blogging systems such as Twitter call for real-time keyword-based querying capabilities.

The workshop provides a forum for discussions of opportunities and challenges in keyword-based search on (semi-)structured data, spatial data, and Web data, and for the presentation of novel techniques in this area.

11th International ACM Workshop on Data Engineering for Wireless and Mobile Access (MobiDE 2012)
http://pages.cs.brandeis.edu/~mobide12/
May 20, 2012
Location: Arizona Ballroom I

MobiDE 2012 is the 11th in a successful series of workshops that aims to act as a bridge between the data management, wireless networking, and mobile computing communities. The workshop serves as a forum for researchers and technologists to discuss the state-of-the-art, present their contributions, and set future directions in relation to data management for mobile and wireless access. MobiDE has been co-located with the annual SIGMOD conference since 2005. MobiDE 2009 marked the 10-year anniversary of the workshop.

1st International Workshop on Scalable Workflow Enactment Engines and Technologies (SWEET 2012)
http://sites.google.com/site/sweetworkshop2012/
May 20, 2012
Location: Arizona Ballroom V

It is a goal to develop systems that are easy to use and understand, but at the same time offer suers put great computational power. The cloud computing model has the potential for advancing this goal in the area of business and scientific data processing. More specifically, cloud computing is facilitating the convergence of workflow-based processing and traditional data management, thereby providing users with the best of both worlds. Recent applications of workflow technology to data-intensive science shows the need for a more robust underlying data management infrastructure. At the same time, workflow-like models and languages are finding their way into data management, a key goal being to make it possible for users with no application development resources, but close to the data domain, to assemble complex data processing pipelines.

The workshop brings together researchers and practitioners to explore the state of the art in workflow-based programming for data-intensive applications and to explore the potential of cloud-based computing in this area.

4th International Workshop on Semantic Web Information Management (SWIM 2012)
http://pamir.dia.uniroma3.it:8080/SWIM2012/Home.html
May 20, 2012
Location: Arizona Ballroom III

As the 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 workshop 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.

15th International Workshop on Web and Databases (WebDB 2012)
http://db.disi.unitn.eu/pages/WebDB2012/
May 20, 2012
Location: Arizona Ballroom IV

WebDB 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. This year's WebDB continues in the spirit of previous years.

8th International Workshop on Data Management on New Hardware (DaMoN 2012)
http://fusion.hpl.hp.com/damon2012/
May 21, 2012
Location: Arizona Ballroom I-III

The workshop brings together researchers who are interested in optimizing database performance on modern computing infrastructure by designing new data management techniques and tools.

As a result of the continued evolution of computing hardware and infrastructure, traditional database architectures that focus solely on I/O optimization increasingly fail to utilize hardware resources efficiently. Multi-core CPUs, GPUs, new memory and storage technologies (such as flash and phase change memory), and low-power hardware impose a great challenge to optimizing database performance. Consequently, exploiting the characteristics of modern hardware has become an important topic of database systems research.

An important goal is to make database systems adapt automatically to hardware characteristics, thus maximizing performance transparently to applications. This calls for interdisciplinary collaboration among database, 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.

5th International Workshop on Testing Database Systems (DBTest 2012)

http://dbtest2012.comp.polyu.edu.hk/index.htm
May 21, 2012
Location: Arizona Ballroom IV

DBTest 2012 brings together researchers and practitioners from academia and industry to discuss key problems and ideas related to testing database systems and applications. The background is that the functionality provided by modern data management systems continues to expand. New applications and usage patterns, the evolution of the underlying hardware and software infrastructure, and increased competition drive the continuous innovation and expansion of these systems. As a result, it has become increasingly expensive to test and tune them, and these stages tend to dominate the release cycle. The workshop focuses on techniques for measuring important properties of database systems, including performance, reliability, security, availability, and robustness. Database systems are defined broadly to include any system that, like a relational database management system, must manage a subtantial amount of data on behalf of applications.

SOCIAL EVENTS

SUNDAY

PODS Reception
Sunday, May 20, 18:00-20:00
Hyatt Regency Arizona Ballroom I–V and South Foyer / Desert Garden

MONDAY

SIGMOD Reception
- **Undergraduate Research Poster Competition**
- **Informal Sponsor/Student Event**

Monday, May 21, 19:00-21:30
Hyatt Regency Terrace Court

TUESDAY

Microsoft Reception
Tuesday, May 22, 21:00-
Hyatt Regency Arizona Ballroom V–VIII

Terrace Court

WEDNESDAY

Conference Banquet
Wednesday, May 23, 18:30-22:30
(Buses start departing at 18:00)
Desert Foothills

Desert Foothills

Spend the evening relaxing in the Sonoran desert. Watch the sunset and then enjoy dinner under the stars. Weather permitting, a tethered hot air balloon will illuminate the evening's festivities. Dress is casual and comfortable shoes are recommended.

LOCAL INFORMATION

In the heart of the Sonoran Desert, Scottsdale welcomes you with the energy of a sophisticated city on the move, the charm and hospitality of days gone by, and a breathtaking landscape that will inspire you to vacation-induced bliss. Amnog others, the city of Scottsdale has been declared:

- 6th most popular travel spot in the US (Travel Leaders, 2010),
- 4th best place to play golf in the world (Golf World, 2008),
- top-25 art destination in the US (American Style, 2010),
- top-100 best places to live (CNN Money, 2010) (RelocateAmerica, 2010),
- top-50 bike-friendly community (Bicycling Magazine, 2010),
- top-100 best communities for young people (America's promise, 2010),
- best place to raise kids (Sunset, 2011),
- best city for babies (Parents, 2010), and
- top-10 city for families (Parents, 2010).

Points of Interest in the Phoenix Metropolitan Area

- **Desert Botanical Garden.** For more than 70 years, the Desert Botanical Garden has been teaching and inspiring visitors from the local community and around the world, providing research, exhibits and more designed to help us understand, protect and preserve the desert's natural beauty. Today, the garden features 50,000+ plant displays showcased in beautiful outdoor exhibits. With approximately fifty acres under cultivation, something is always blooming at the Desert Botanical Garden.
- **Phoenix Zoo.** Voted one of the nation's top 5 zoos for kids, the Phoenix Zoo is home to more than 1,300 animals including 200 endangered species and participates in 37 Species Survival Plans (SSPs) to breed and maintain endangered species populations in captivity.

- **Phoenix Symphony.** Based in Phoenix Symphony Hall (opened in 1972, renovated in 2005, and seating 2,312), the 76-member ensemble is Arizona's only full-time, professional orchestra. The orchestra offers 275 concerts and presentations in an annual season running from September to May. Events in the second half of May 2012 include Concierto de Aranjuez and Cirque de la Symphonie.
- **Arizona Science Center.** Arizona Science Center offers educational fun for the whole family, from pre-K to gray! From family-friendly presentations to hands-on gallery programs, get your hands on science!
- **Arizona Diamondbacks.** In their first five seasons, the Diamondbacks won three division titles, a National League pennant, and a World Series championship in 2001. You can check them out at the Chase Field located in downtown Phoenix.
- **Heard Museum.** The Heard Museum actively collects American Indian fine art, and a variety of paintings, drawings, prints, photography and sculpture can be found throughout the museum's galleries and grounds. Dedicated to the sensitive and accurate portrayal of Native arts and cultures, the Heard is an institution that successfully combines the stories of American Indian people from a personal perspective with the beauty of art. Partnerships with American Indian artists and tribal communities provide visitors with a distinctive perspective about the art and cultures of Native people, especially those from the Southwest.
- **Taliesin West Frank.** Lloyd Wright began building this desert masterpiece in 1937 as his personal winter home, studio, and architectural campus. Located on the beautiful Sonoran desert in the foothills of the McDowell Mountains in northeast Scottsdale, the site offers a broad range of guided public tours. Visitors experience firsthand Wright's brilliant ability to integrate indoor and outdoor spaces.
- **Lost Dutchman State Park.** Superstition Mountain and the Dutchman's Lost Mine are synonymous with Arizona lost mine lore. Lost Dutchman State Park is home to the Lost Dutchman Museum and the Goldfield Ghost Town.
- **Arcosanti.** An experimental town in the desert of Arizona, built to embody Paolo Soleri's concept of arcology - the fusion of architecture with ecology. Arcosanti is an urban laboratory focused on pursuing lean alternatives to urban sprawl through innovative design with environmental accountability.
- **Old Town Scottsdale.** Take a walk through the historic Old Town Scottsdale and get a flavor of the old west. Wooden sidewalks and rustic traditions are preserved as you tie your horse to each shop's hitching post. Browse the many quaint stores that line the streets of Old Town and you'll find a mix of traditional Western-wear retailers, Southwestern specialty shops and tasty treats.
- **Greasewood Flat.** The original stagecoach stop between Fort McDowell and Phoenix, the Greasewood Flat bar is housed in a 120+-year-old bunkhouse. Greasewood Flat (which serves hearty burgers and some wicked libations) and the nearby Reata Pass Steakhouse are located at an elevation 2000 feet above the valley floorm making the night air 10 degrees cooler.

In addition, many parks have been established to preserve the desert landscape in areas that would otherwise quickly be developed with commercial and residential zoning. Aside from the Desert Botanical Garden listed above, the most noteworthy park is South Mountain Park, the world's largest municipal park with 16,500 acres (67 km^2); others include Camelback Mountain, Sunnyslope Mountain, also known as "S" Mountain, and Piestewa Peak, which boasts dozens of miles of trails to enjoy the glory of the Sonoran Desert in relative solitude. Encanto Park, named one of the "Top 12 Best City Parks in America!" by the Forbes Magazine, is the city's largest and primary urban park, and lies just northwest of downtown Phoenix. Papago Park in east Phoenix is home to both the Desert Botanical Garden and the Phoenix Zoo, as well as a few golf courses.

Outdoor Adventures

Explore Arizona in a variety of ways and locations - in the Phoenix area, the enchanting red rocks of Sedona, or the majestic Grand Canyon. The sky's the limit - literally! Adventures include hot air balloon rides, hiking, scenic tours, and more. You may check the 360 Adventures Web site for more details on some exciting tours.

- Tours available in the Phoenix metropolitan area in the morning and afternoon for a duration of about 2 hours include hot air balloon rides, horseback riding, mountain biking, hiking, kayaking, and Hummer and ATV tours.
- Sedona tours depart in the morning for a duration of about 10 hours and include scenic, hiking and mountain biking.
- Grand Canyon tours depart in the morning for a duration of about 12 hours and include scenic and hiking tours.
- Explore Lake Powell and Antelope Canyon at Page if you have two or three days, which offer unique scenery and a most-photographed spot in the American Southwest.
- Enjoy the Petrified Forest National Park for globally significant Late Triassic fossils and Painted Desert if you have two or three days.
- Admire the gigantic saguaros at the Saguaro National Park at Tucson for a day trip.

360 Adventures offers some of the above tours and requires online scheduling of your outdoor adventures 7 days in advance for availability. Mention the ACM SIGMOD conference on your reservation under Special Requests to be grouped with conference attendees when possible and to get a potential group discount. (Some tour vendors are providing discounts for conference attendees in groups of 10 or more.) Also note on the online reservation whether you want a morning or afternoon time for the local tours. After submission of the online reservation, you must call 360 Adventures to secure the reservation with a credit card, which will be charged 72 hours before the adventure per the cancellation policy. 360 Adventures is also on site at the conference hotel - Hyatt Regency Scottsdale Resort and Spa at Gainey Ranch.

Resort Restaurants

Surround yourself with the sights and sounds of the Sonoran Desert as you dine al fresco in our open-air Scottsdale restaurants. Wood-fired specialties abound at SWB, our Southwest bistro, offering regional favorites such as scrambled eggs and chorizo breakfasts or seared chili marinated shrimp and scallop dinners. Alto ristorante e bar offers flavorful Italian fare with an accent on "fresh seasonal" ingredients. Dinner in Alto comes complete with a complimentary gondola ride through the resort's waterways. When just a light bite is in order, be sure to stop by Canyon Market or Water Garden, our poolside cafe. When evening descends, settle in to a comfortable armchair in our open-air Lobby Bar and sip your favorite beverage while enjoying live entertainment and the company of friends and fellow guests.

SWB, a Southwest bistro

Relax and enjoy a delectable breakfast, lunch, dinner and Sunday brunch indoors or out, set against a backdrop of the majestic McDowell Mountains. Begin your day on the sun-drenched patio of our casual Southwestern Scottsdale restaurant with a steaming cup of coffee and our famous Sticky Bun French Toast. Watch our chefs create wood-fired specialties with a distinctly Southwest flair featuring fresh seasonal ingredients.

Hours: Breakfast: 6:30a.m. - 11:30p.m., daily
Lunch: 11:30a.m. - 2:30p.m., daily
Dinner: 5:00p.m. - 10:30p.m, daily
Sunday Brunch: 6:30a.m. - 12:00p.m

Reservations**:** Please call 480-444-1234 ext. 79, or visit http://www.opentable.com/.
Phone: 480-444-1234 ext. 79

"Wine Me, Dine Me" at SWB

Take a journey deep into the heart of authentic Southwestern flavor as Chef Juan and his culinary staff prepare a gourmet repast for you. A specially crafted four-course menu complete with tequila and wine pairings will imprint an indelible memory of your Sonoran Desert experience.

Reservations recommended. Please call 480-444-1234 ext. 79
Hours: Tuesday through Saturday: 5:00p.m. - 9:00p.m

Alto Ristorante e Bar

Add a romantic element to your dining experience at our enchanting Scottsdale Arizona restaurant. Cross over the bridge and through a courtyard to find the vibrant, convivial setting of Alto ristorante e bar. This comfortable Italian restaurant offers fresh and flavorful dishes with an emphasis on “the season.” Formaggio & Salumi Cellar, Alto’s unique collection of Italy’s best artisanal cheeses and meats are complimented by their extensive selection of wines from around the world (with an emphasis on Italy). Dinner is served both indoors and alfresco nestled around terrace fire pits, creating a convivial atmosphere. Top off your evening with a complimentary gondola ride along the resort's waterways, complete with singing gondoliers. Complimentary Gondola Boat Ride with the purchase of a dinner entree.

Attire: Resort Casual
Hours: 5:30p.m. - 10:30p.m., daily
Reservations: Please call 480-444-1234 ext. 79, or visit Alto - OpenTable
Phone: 480-444-1234 ext. 79

Noh

Sushi, sashimi and traditional Japanese dishes with a contemporary twist.

Hours: 5:00p.m. - 10:00p.m., Thursday through Saturday

Water Garden

Dine poolside. Have your favorite foods and refreshments delivered right to you so you don’t need to leave the fun of our amazing 2.5 acre water playground. The Water Garden offers a wide variety of lunchtime favorites, tableside or poolside, including paninis, sandwiches, salads and refreshing beverages. Drink to your health at our seasonal Waterfall Juice Bar. Order a healthful juice or smoothie prepared just for you featuring the freshest of fruits and vegetables, juiced to order.

Hours: 11:00a.m. - 5:00p.m., daily (Seasonal)

Canyon Market

Canyon Market, our convenient 24 hour market and hiking lifestyle store is the perfect solution to “what do I want?” Take a little late night snack back to your room, satisfy your need for your favorite Starbucks brew or fortify yourself while you’re on the go at Canyon Market. Located in the entry court, our hiking lifestyle store that doubles as a 24 hour market, features Starbuck's coffee, ice cream, sandwiches, salads and daily organic offerings along with performance food and gear essential for an active lifestyle.

Hours: Open 24 hours, daily

Restaurants at Walking Distance

Chez Vous Crepes & Gelato

This cozy and authentic Parisian creperie features savory and sweet crepes to breakfast crepes, fresh croissants, salads, triple-layer French sandwiches and made-to-order baguettes, all freshly prepared and paired with ten natural and organic gelato flavors. Chez Vous is open Monday through Saturday from 7:30 am to 4:00 pm for breakfast and lunch, remaining open Monday-Thursday evenings until 9:00 pm, Friday-Saturday evenings until 10:00 pm for gelato and beverages exclusively.

Mon-Sat 7:30am-4pm (breakfast & lunch), Mon-Thur 7:30am-9:00pm, Fri-Sat 7:30am-10pm
Phone: (480) 443-2575

McCormick & Schmick's

McCormick & Schmick's offers guests dozens of fresh seafood varieties at the peak of their seasons. Our menus are printed daily featuring a "Fresh List" of more than 30 seafood species. We're also famous for our happy hour (Monday through Friday from 3-6 p.m. and 9-11 p.m.) during which we offer a .95 bar menu. Please visit our web site to make reservations online, or for more information.

Mon-Fri: Lunch 11-4, Sat-Sun: Lunch 11:30-4, Sun-Thurs: Dinner 4-10, Fri-Sat: Dinner 4-11
Phone: (480) 998-2026

Paradise Bakery & Cafe

Fresh sandwiches, salads made to order and fresh cookies and muffins.

Mon-Sat 6am-8pm, Sun 6am-6pm
Phone: (480) 951-2500

Pei Wei

From the exotic aromas rising from the sizzling woks, to the signature red floors and casually hip decor, Pei Wei (pronounced Pay Way) combines distinct Asian dishes with the American lifestyle. Dine In or Take Away.

Sun-Th 11am-9pm, Fri-Sat 11am-9:30pm
Phone: (480) 365-6000

SOI Four Bangkok Eatery

From the Siri Group, a family business, which owns and operates four highly acclaimed restaurants in the San Francisco Bay Area: the original Sweet Basil Thai Restaurant, Basil Tai Restaurant and Bar, Basil Canteen and Soi Four Bangkok Eatery, is proud to bring our newest Soi Four Bangkok eatery to Scottsdale. “Soi” is a Thai word for side streets or alleyways that, together, form the essential lifeline that makes up the fabric of Bangkok city. Like the city, our menu is a modern take on the traditional through the use of fresh local ingredients, sustainable practice and creative approach to authentic Thai dishes. At Soi 4, we also offer Thai inspired artisanal cocktails and wide varety of local and imported beers… influenced, accented, tinged anything but “fusion”. Chi yo – Cheers!

Mon-Sat 11:30am-2pm, Mon-Thur 4:30pm-9:00pm, Fri-Sat 4:30pm-10pm
Phone: (480) 778-1999

The Coffee Bean & Tea Leaf

Founded in 1963, The Coffee Bean & Tea Leaf brings great coffees of the world and exotic, rare teas to America through our shops.

M-F 5:30am-8pm, Sat.-Sun. 6am-8pm
Phone: (480) 315-9335

Village Tavern

The Village Tavern offers a favorable combination of the best of the best traditional fare with newer and more modern menu items.

Mon-Th 11am-Midnite, Fri-Sat 11am-1am, Sun 11am-10pm
Phone: (480) 951-6445

Wally's American Pub N' Grille Restaurant Bar

Wally's American Pub N' Grille restaurant and bar are a local favorite serving great comfort food in the Scottsdale area!
Mon-Sun 11:00 am-9pm

Other Restaurants

Our Favorites

Olive and Ivy *(480) 751-2200* *7135 E. Camelback Rd.* American Mediterannean cuisine at the waterfront.	**Cowboy Ciao** *(480) 946-3111* *7133 E. Stetson Dr.* Postmodern Western bistro featuring a unique american fare menu.	**Elements at the Sanctuary Resort** *(480) 607-2300* *5700 E. Mcdonald Dr.* A monthly-changing menu of American cuisine, Asian elements, and local organic foods. Spectacular mountainside dining room.	**T'Cook's at the Royal Palms Resort** *(602) 808-0766* *5200 E. Camelback Rd.* Serving Mediterranean fare and specialites from the rotisserie roasting fireplace in an elegant setting.
Lon's at the Hermosa Inn *(602) 955-7878* *5532 N. Palo Christi Rd.* Contemporary menu with shades of French and Italian.	**Modern Steak** *(480) 423-7000* *7014 E. Camelback Rd.* Steak, appetizers, prime-aged beef, seafood with a unique "modern" spin.		

Asian

Flo's *(480) 661-8883* *14850 N. Frank Lloyd Wright Blvd.* Unique Chinese cuisine.	**Malee's** *(480) 947-6042* *7131 Main St.* Traditional Thai cuisine.	**Roka Akor** *(480) 306-8800* *7299 N. Scottsdale Rd.* Artistic little tastes come from the robata (charcoal) grill or the shushi/shashimi bar.	**Sapporo** *(480) 607-1114* *14344 N. Scottsdale Rd.* Japanese, Chinese, Thai, and Vietnamese cuisine, plus sushi and teppanyaki.
Stingray Sushi *(480) 941-4460* *4302 N. Scottsdale Rd.* Modern interiors and some of the most creative sushi the desert has ever seen.	**Sushi Roku** *(480) 970-2121* *7277 E. Camelback Rd.*	**Christopher's & Crush Lounge** *(602) 522-2344* *2502 E. Camelback Rd.* Casual atmosphere, wines matched to entrees.	**Metro Brassiere** *(480) 994-3663* *7114 E. Stetson Dr.* Rustic French cooking with influences from the American South.
Zinc Bistro *(480) 603-0922* *15034 N. Scottsdale Rd.* Parisian-styled bistro with a extensive wine list.			

Italian

Arrivederci *(480) 922-8225* *7101 E. Thunderbird Rd.*	**Maggiano's** *(480) 333-4100* *16405 N. Scottsdale Rd.*	**North** *(480) 948-2055* *15024 N. Scottsdale Rd.* New age Italian at Kierland Commons.	**Un Bacio** *(480) 609-6969* *7704 E. Doubletree Ranch Rd.*
Veneto Trattoria *(480) 948-9928* *6137 N. Scottsdale Rd.* Specializes in fine, classic Italian fare.			

Indian

Jewel of the Crown *(602)-840-2412* *7373 Scottsdale Mall*			

Mexican

Blanco Tacos + Tequila *480-305-6692* *6166 N. Scottsdale Rd.* Mexican cuisine	**La Hacienda** *(480) 585-4848* *7575 E. Princess Dr.*	**Los Olivos** *7328 2nd St.* *(480) 946-2256* *15544 N. Pima Rd.* *(480) 596-9787* Family owned restraunt serving Sonoran Mexican.	**Los Sombreros** *(480) 994-1799* *2534 N. Scottsdale Rd.* Creative dishes in a cozy atmoshphere.

Seafood

Chart House *(480) 951-2550* *7255 McCormick Pkwy.* Serves fresh seafood, chicken, steaks, and pasta in a casual setting with a view.	**Eddie V's** *(480) 538-8468* *20715 N. Pima Rd.* Fresh seafood served in a bistro setting with live music at the lounge.	**Ocean Club** *(480) 443-8555* *15045 N. Kierland Blvd.* Fine dining seafood restaurant with music in the lounge.	**Wildfish** *(480) 994-4040* *7135 E. Camelback Rd.* Fresh seafood and fun atmosphere at the waterfront.

Spanish

Deseo *(480) 624-1015* *6902 E. Greenway Pkwy.* Spain's best at the Westin Kierland.			

Southwestern

The Mission *(480) 636-5005* *3815 N. Brown Ave.* Zesty plates and grilled entrees set in an eclectically elegant atmosphere.	**Roaring Fork** *(480) 947-0795* *4800 N. Scottdale Rd.* Lively bistro featuring Chef Mcgrath's Western creations.	**Vincent on Camelback** *(602) 224-0225* *3930 E. Camelback Rd.* Upscale dining room featuring French cooking techniques with Southwestern ingredients. * jacket suggested	**Z'Tejas Grill** *7014 E. Camelback Rd.* *(480) 946-4171* *10625 N. Tatum Blvd.* *(480) 948-9010* Casual Southwest grill featuring unique appetizers, seafood favorites, and grilled meats.

Steakhouse

Bourbon Steak *(480) 513-6002* *7575 E. Princess Dr.* Elegant steakhouse offering finest cuts of beef, poultry, and seafood.	**Capital Grill** *(480) 348-1700* *16489 N. Scottsdale Rd.* Steak and seafood.	**Don & Charlies** *(480) 990-0900* *7501 E. Camelback Rd.* American rib and chop house featuring prime aged steaks, savory ribs, and barbecued chicken.	**Fleming's** *(480) 596-8265* *6333 N. Scottsdale Rd.* Prime steakhouse and wine bar.
Fogo de Chao *(480) 609-8866* *6300 N. Scottsdale Rd.* Brazilian steakhouse.	**Mastro's** *(480) 585-9500* *8852 E. Pinnacle Peak Rd.* Signature prime steaks featuring an upscale dining room with live music in the lounge.	**Mastro's City Hall** *(480) 941-4700* *6991 E. Camelback Rd.* Elegant setting featuring live jazz music.	

Breakfast

Breakfast Club *(480) 222-2582* *4400 E. Stetson Dr.* Specializing in the morning meal, relaxing atmosphere.	**Butterfields Pancake House** *(480) 951-6002* *7388 E. Shea Blvd.*	**Daily Dose** *(480) 994-3673* *4020 N. Scottsdale Rd.* Hearty breakfasts and all-day fare	**The Good Egg** *(480) 991-5416* *6149 N. Scottsdale Rd.*

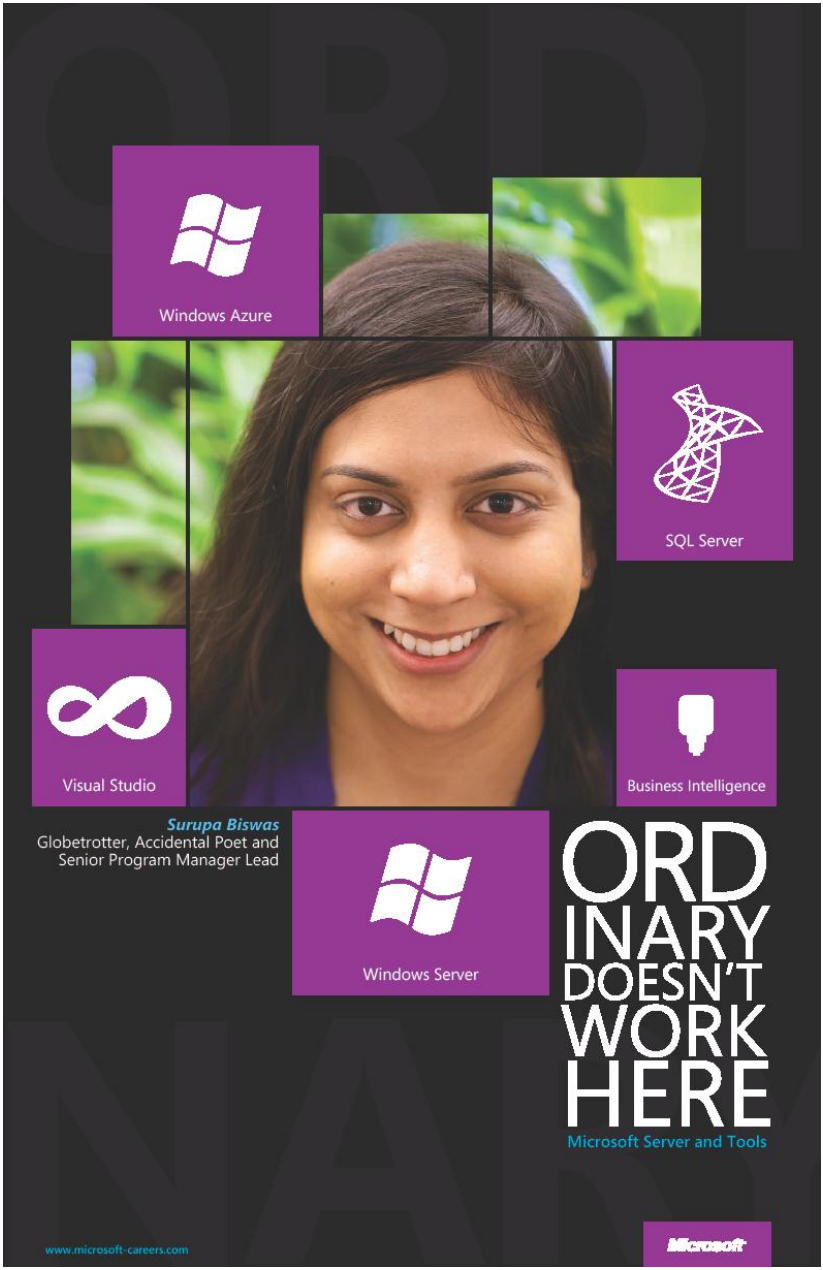
Windows Azure
SQL Server
Visual Studio
Business Intelligence
Surupa Biswas
Globetrotter, Accidental Poet and
Senior Program Manager Lead
Windows Server
ORD
INARY
DOESN'T
WORK
HERE
Microsoft Server and Tools
www.microsoft-careers.com
Microsoft

IBM's rich history in information management encompasses revolutionary breakthroughs ranging from the invention of the relational database by researchers in the Silicon Valley to advanced text analytics capabilities demonstrated by Watson, which proved its computing prowess on the quiz show *Jeopardy!* earlier this year. Key technologies resulting from IBM research projects such as System R, R*, Garlic, DB2 pureXML, System S and Clio and System T, along with the invention of the leading data mining algorithms, have shaped the industry.

IBM Research is the leading corporate research institution in the world, a global community of 3,000 world-class scientists who have been honored with five Nobel Prizes, six Turing Awards, five National Medals of Science, 80 memberships in the National Academies and 11 inductions into National Inventors Hall of Fame. For the **eighteenth** straight year, IBM received the most U.S. patents of any company worldwide.

IBM works with clients around the world to help their businesses grow "smarter" as the planet becomes more instrumented, interconnected and intelligent. Through its software, hardware, services, consulting and R&D expertise, IBM helps create systems that lead to less traffic, healthier food, cleaner water, and safer cities.

Our Ground Floor is in the Cloud

Do you want to develop and support innovative solutions that deliver the future of IT through cloud computing? As an intern or new college graduate, your innovative ideas matter from day one, so bring them here and work alongside technical gurus who are eager to share with and learn from you too. The audacity to challenge constraints and problem-solve for tomorrow starts today, and it starts with you.

There are many opportunities to get involved and pursue your passions at VMware. The VMware Academic Program (VMAP) is a comprehensive global program designed specifically for the academic community. VMAP offers a wide variety of programs that include research funding, paper publication incentives, academic licensing and PhD recruitment.

To learn more, connect with us:




labs.vmware.com @vmwacademic vmware.com/careers/ur @VMwareU

Contact Decision Makers Directly and Shorten Your Sales Cycle

Sales Navigator includes LinkedIn InMail, the powerful contact tool that lets you send messages to decision makers even if they're not in your network. According to a CSO Insights study, on average 8 InMails can generate 1 new sales opportunity.[1]

Find and Be Found with Trusted Data

Make your LinkedIn profile work harder for you. With your LinkedIn profile, present a professional impression of both you and your company for interested inbound leads. The power of LinkedIn's professional graph and insights extends to prospects too. LinkedIn members self-author their profile information, so you can trust that it's accurate and up to date. In fact, 87% of sales reps found more information about people or companies using Sales Navigator that they wouldn't have found otherwise.[2]

Prioritize Your Accounts and Territory

Don't waste time pursuing ineffective leads. Sales Navigator's Lead Builder tool helps you build a contact list based on criteria such as seniority, title, function, and industry, and prioritizes that list by the number of people you know at that company. And with other LinkedIn tools such as Profile Organizer and Saved Searches, you can track profiles, organize them into folders, and add your own sales notes.

Expand Your Network and Leverage Your Common Connection

With 150 million members, LinkedIn offers real connections to the people and companies you need to reach. While employees with more coworker connections receive more introductions, only 53% of sales professionals have leveraged their coworkers to facilitate introductions to new opportunities.[2] With Team Link from LinkedIn, you automatically see who on your sales team is a first-degree connection with your prospect, resulting in warm introductions and more deals won.

Drive Deeper Adoption of Your Existing CRM Investment

Sales Navigator users using LinkedIn within their CRM view an average of 27 LinkedIn profiles a day, compared to the broader population of sales professionals who view just five profiles per day.[1] By integrating your CRM with social channels like LinkedIn, you are researching and connecting with your prospects in the environment where they live and prospecting smarter within your existing CRM.

MISSION:

Google's mission: to organize the world's information and make it universally accessible and useful.

RESEARCH AND DEVELOPMENT:

Big Challenges, Big Resources, Big Impact

The challenges of realizing this mission, and of operating Google's global computing system, span the breadth of computer science and engineering. The research and development activities that we undertake in the areas of search, advertising, algorithms, social, data management, systems infrastructure, machine learning, programming languages, back our major products. Google has a tremendous number of exciting challenges that only arise through the vast amount of data and sheer scale of systems we build. R&D happens across the company, on many different teams, making the line between research and development at Google wonderfully blurred. We actively engage with academics through various initiatives such as our Research Awards Program, Visiting Faculty Program, and PhD Fellowship Program. We understand the value of a collaborative ecosystem; we launch products that demonstrate progress and functionality, and have a strong culture of releasing code as open source.

We are hiring!

http://www.google.com/intl/en/jobs/students/
http://research.google.com
Add "Research at Google" to your circles on Google+
http://goo.gl/cKrpz

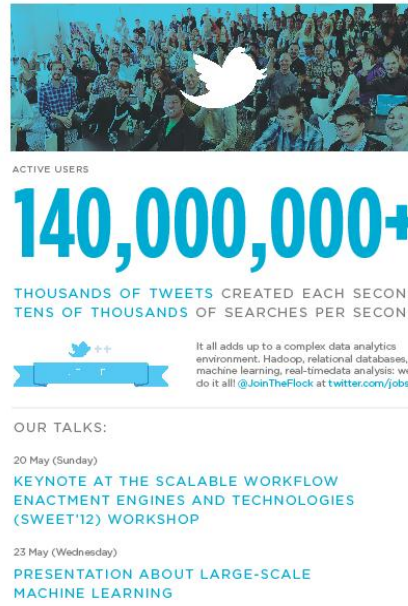

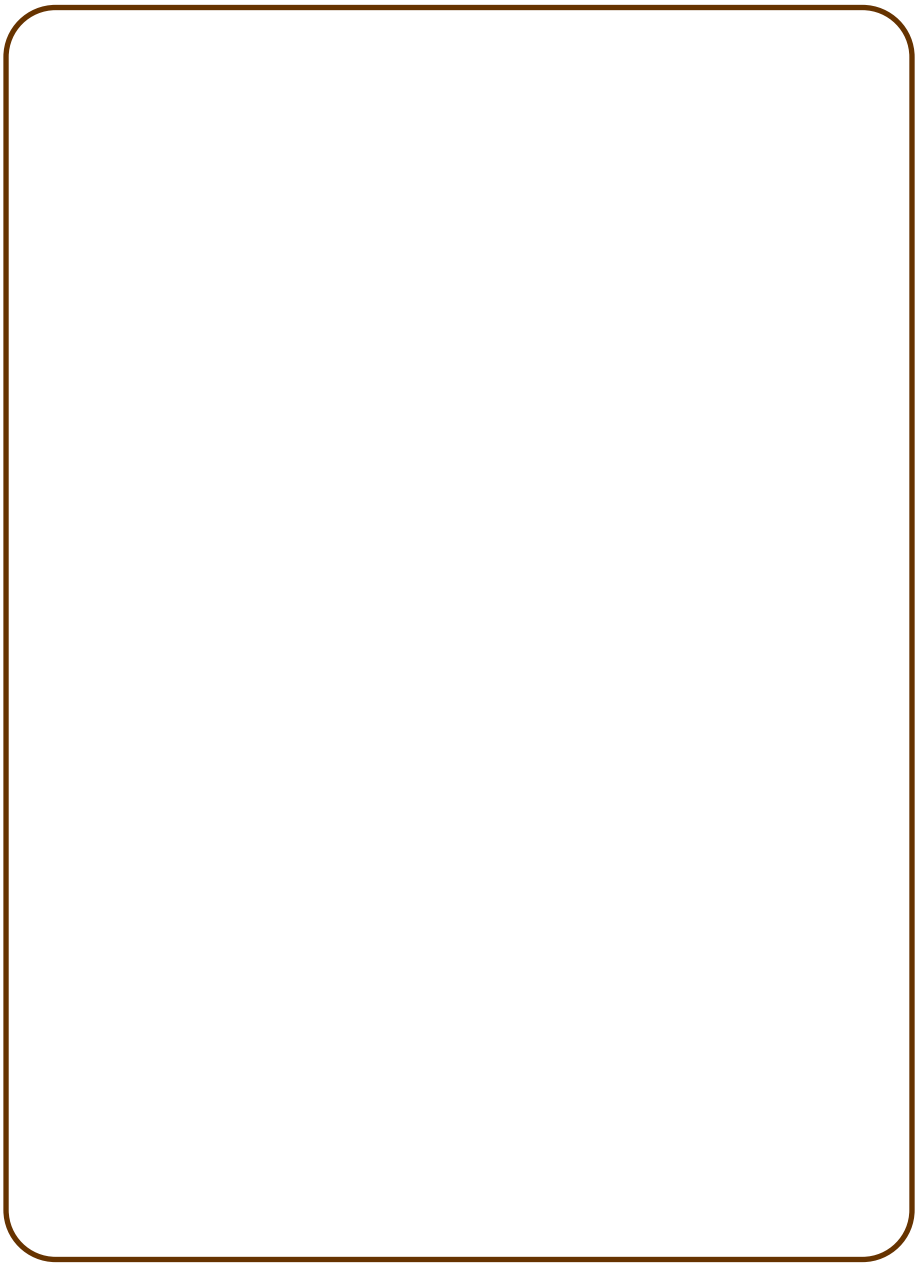

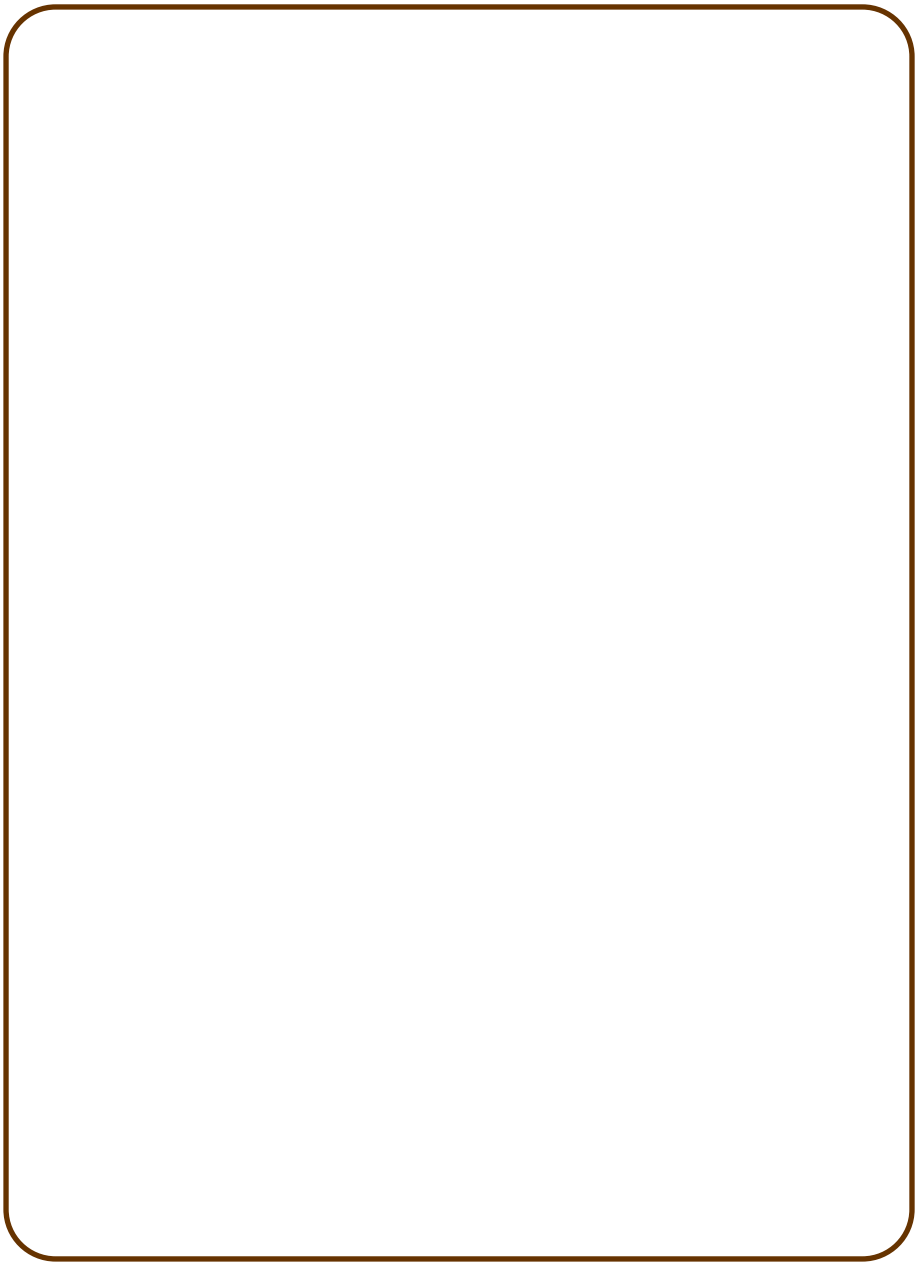

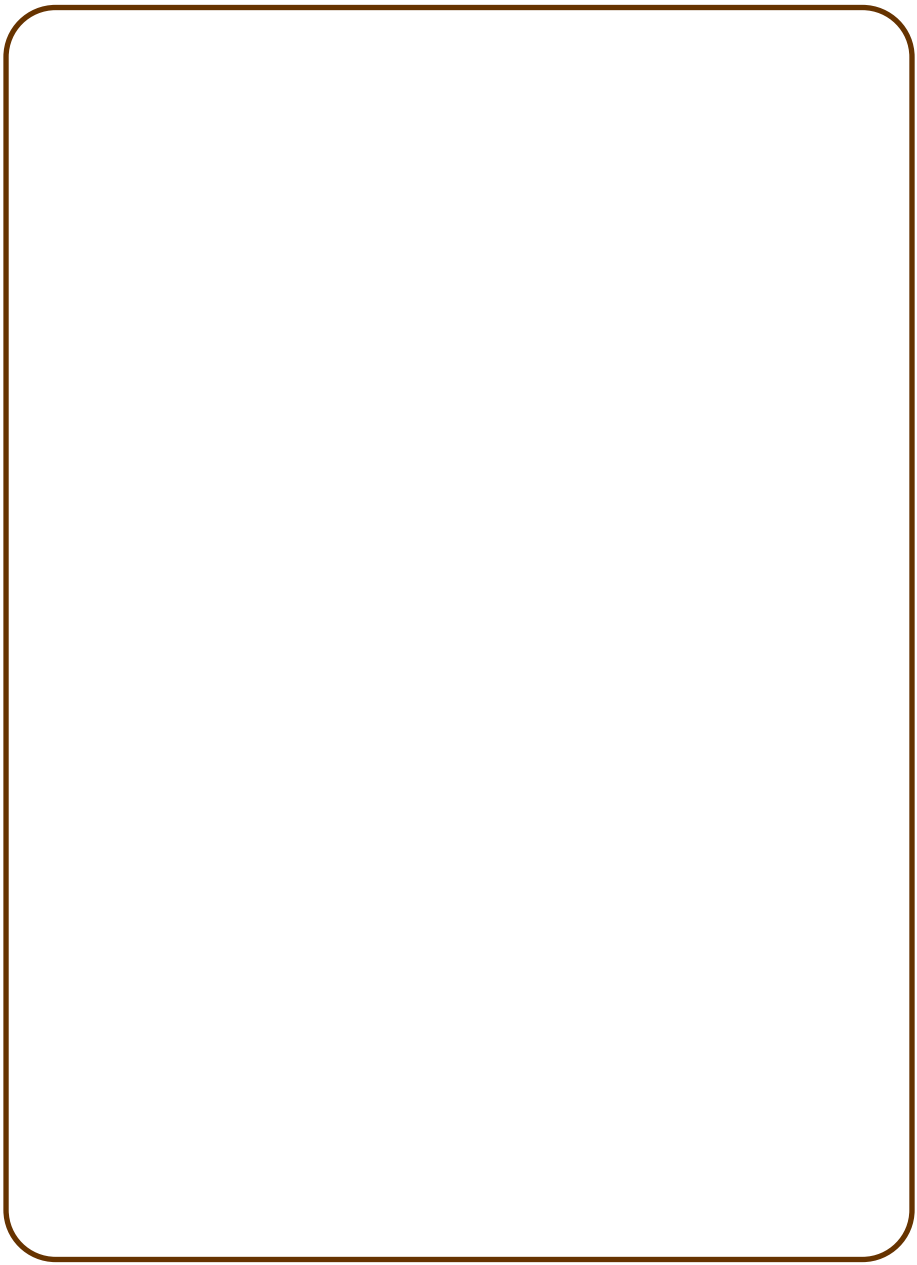

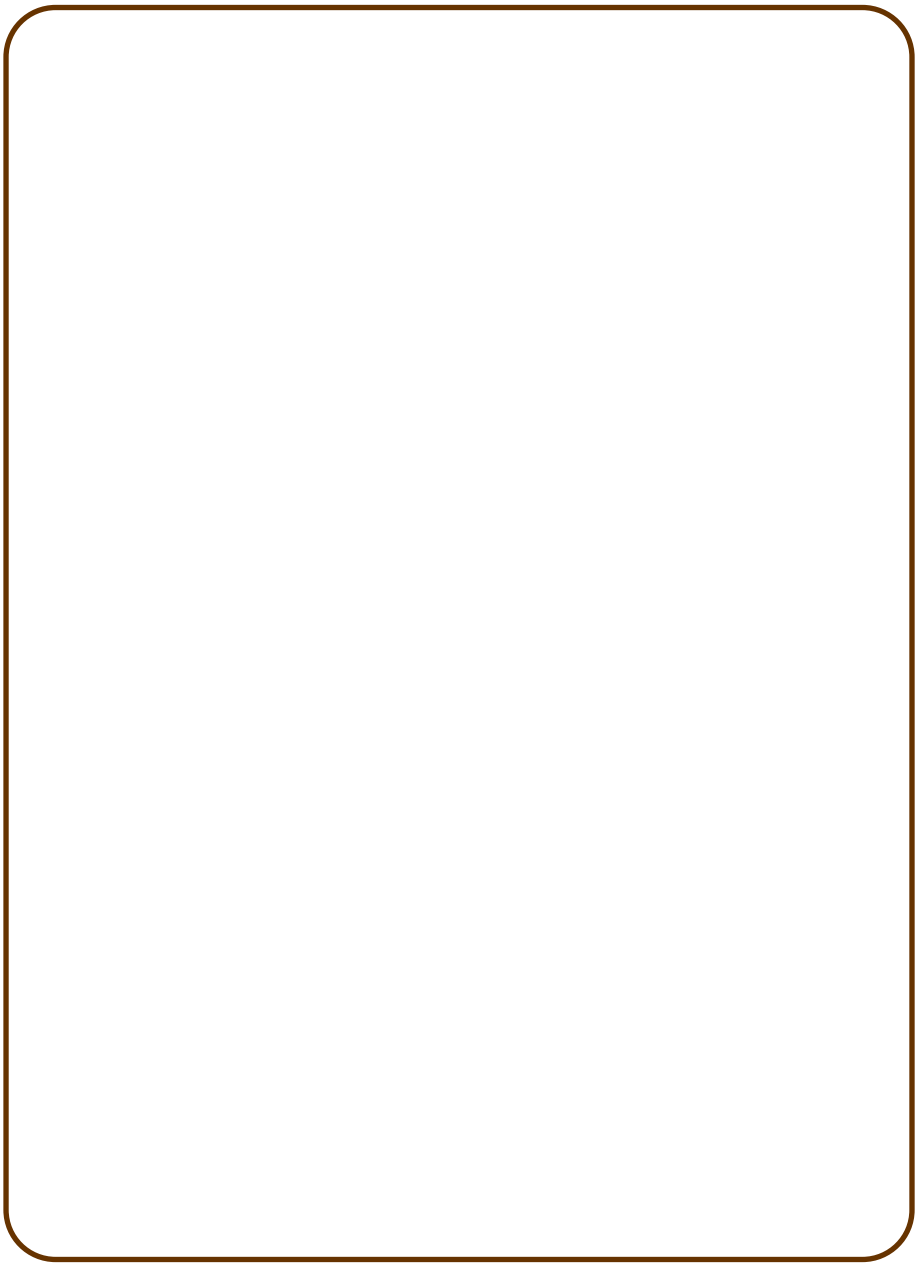

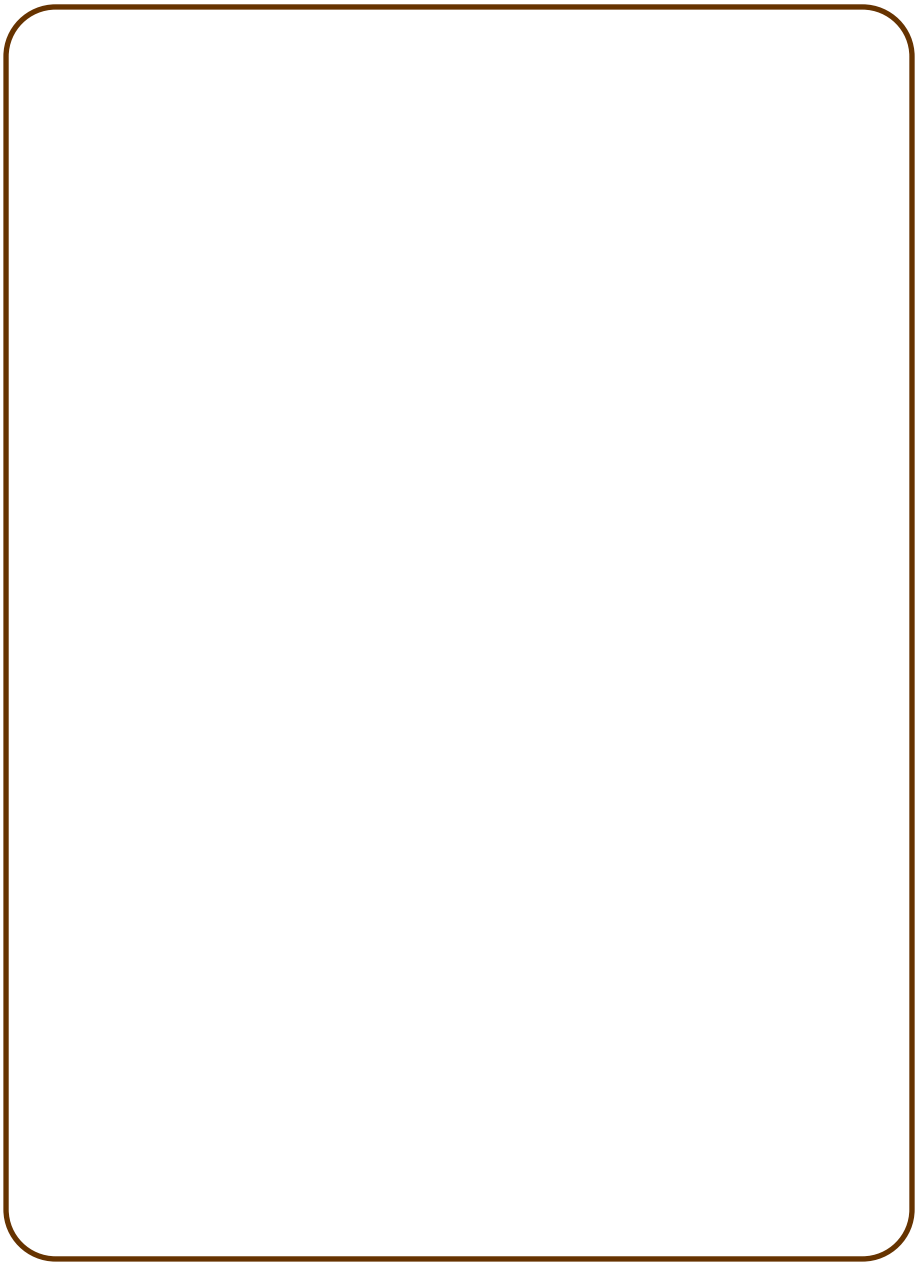

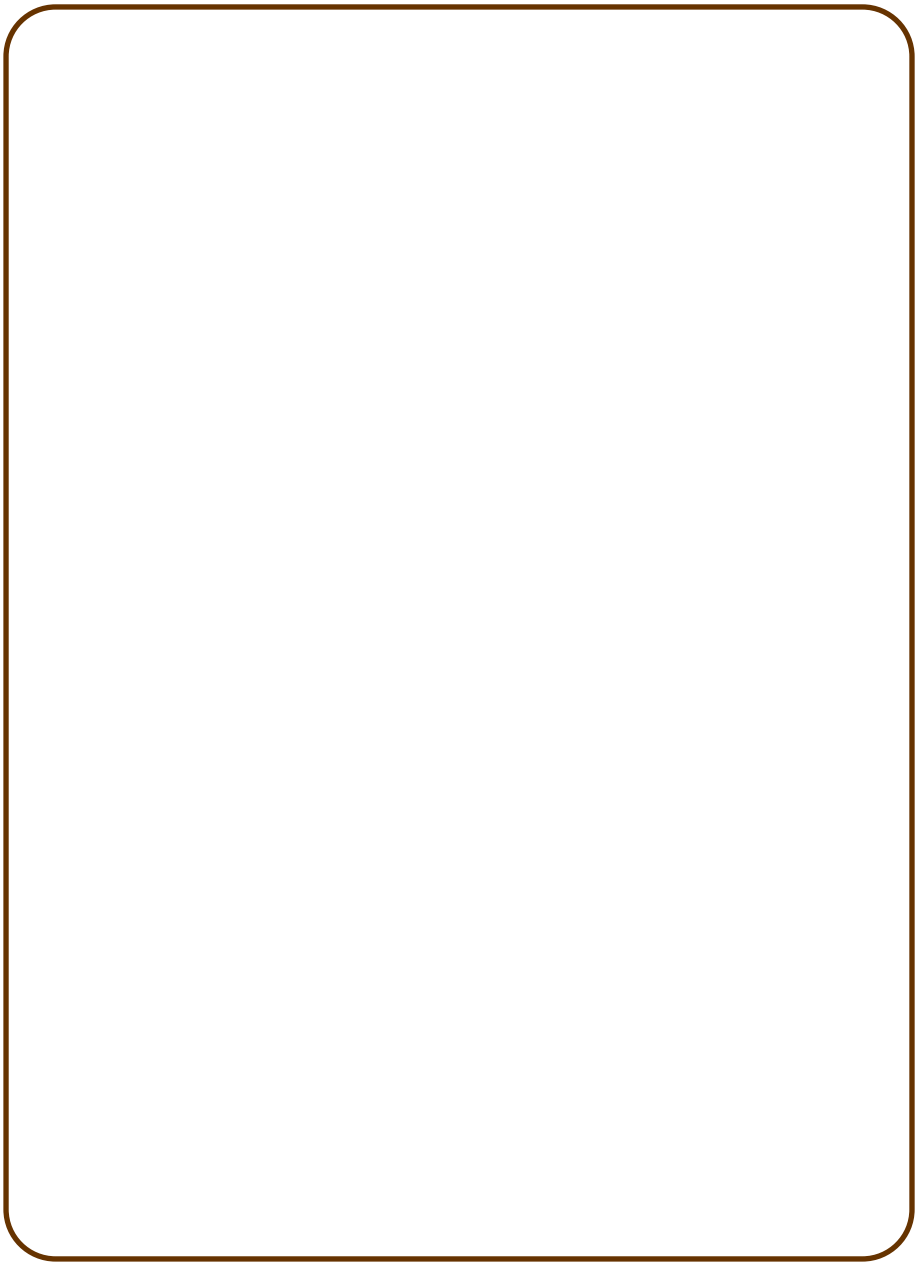

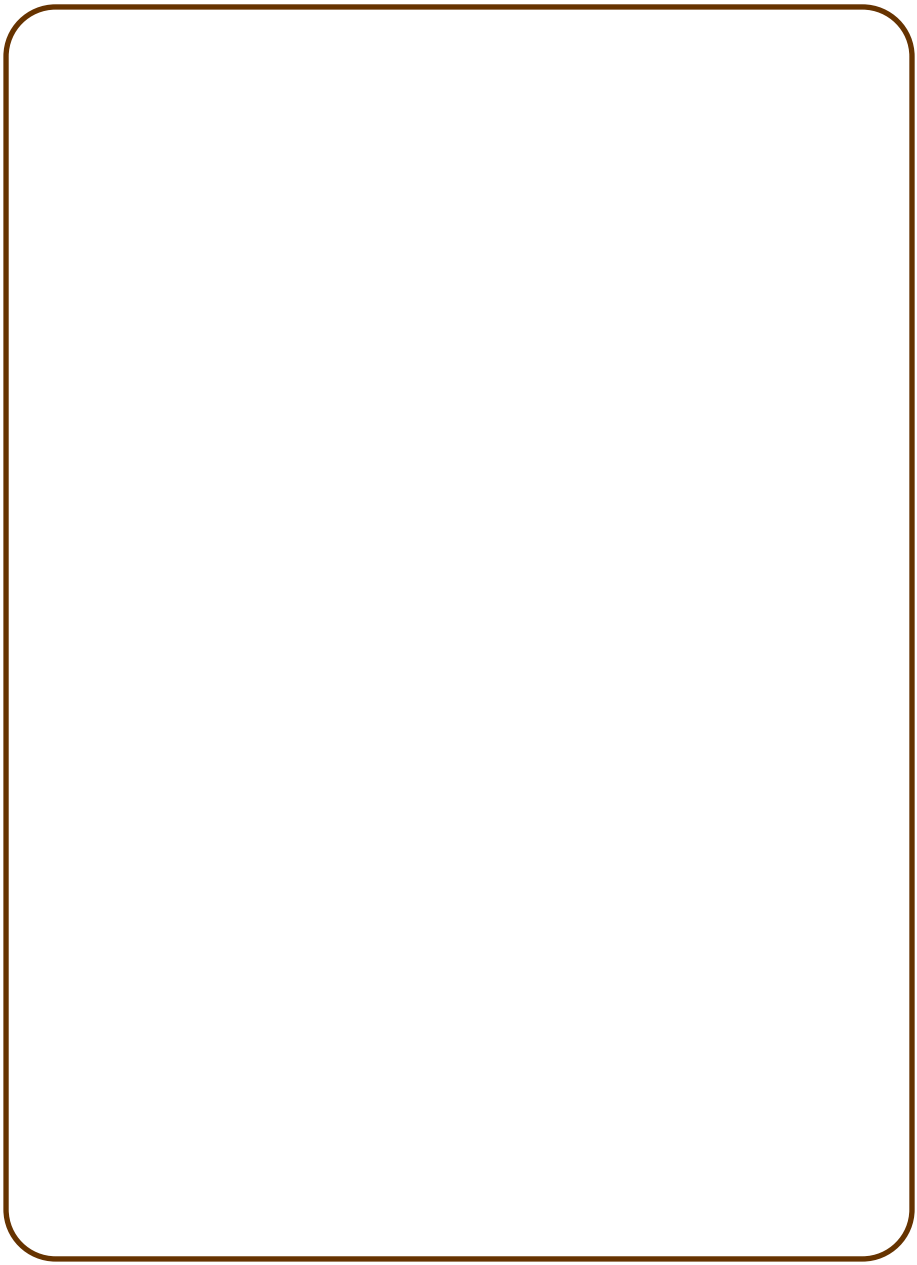

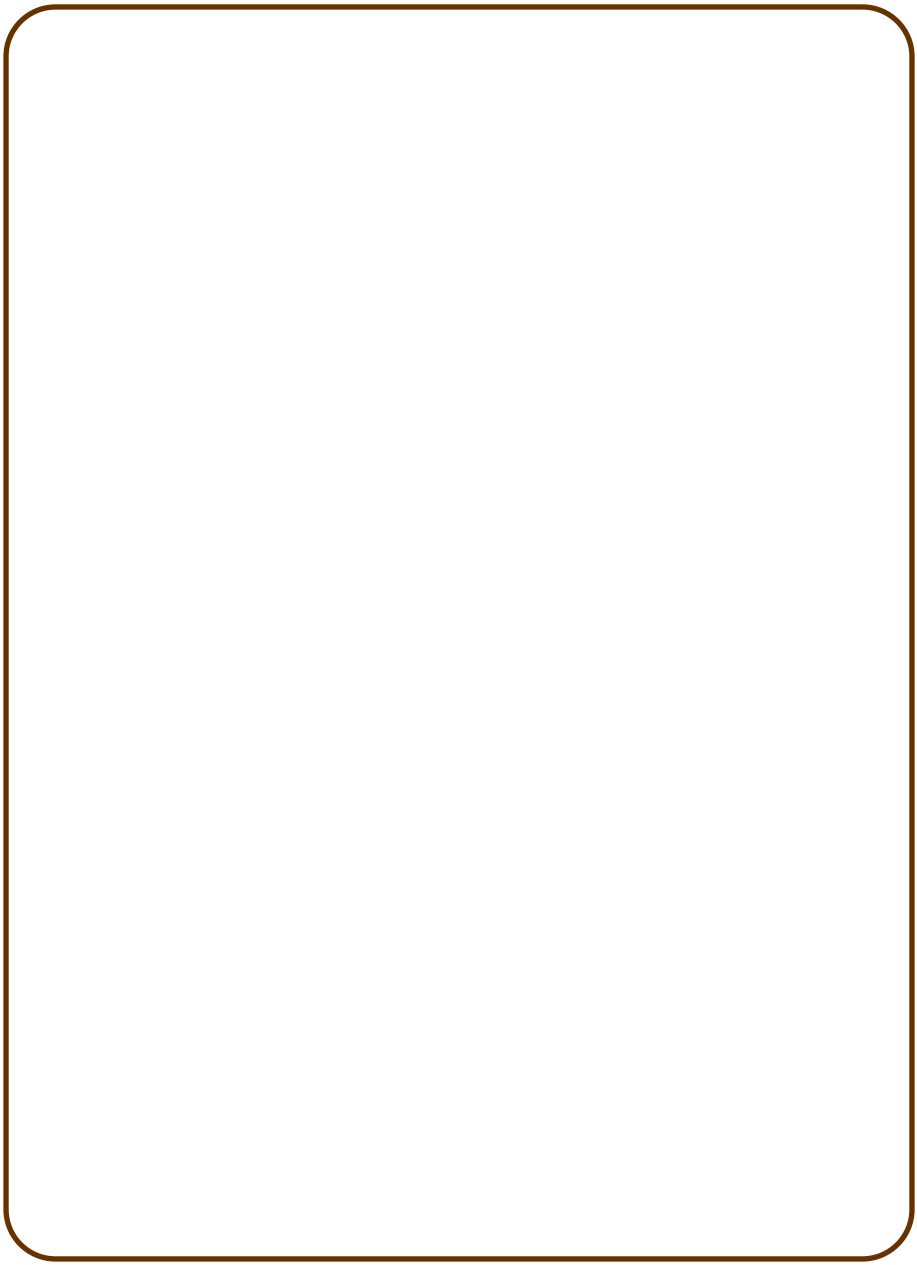

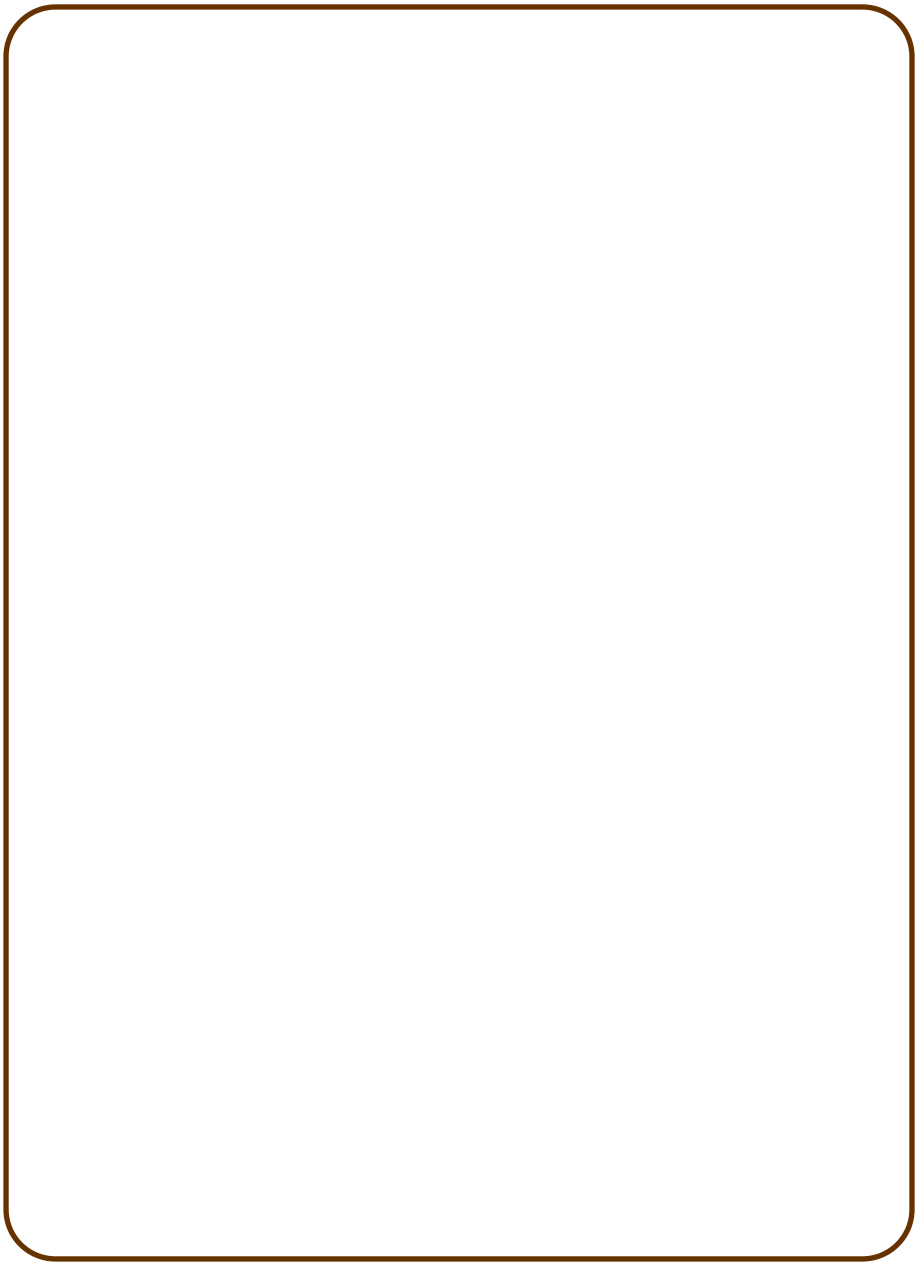

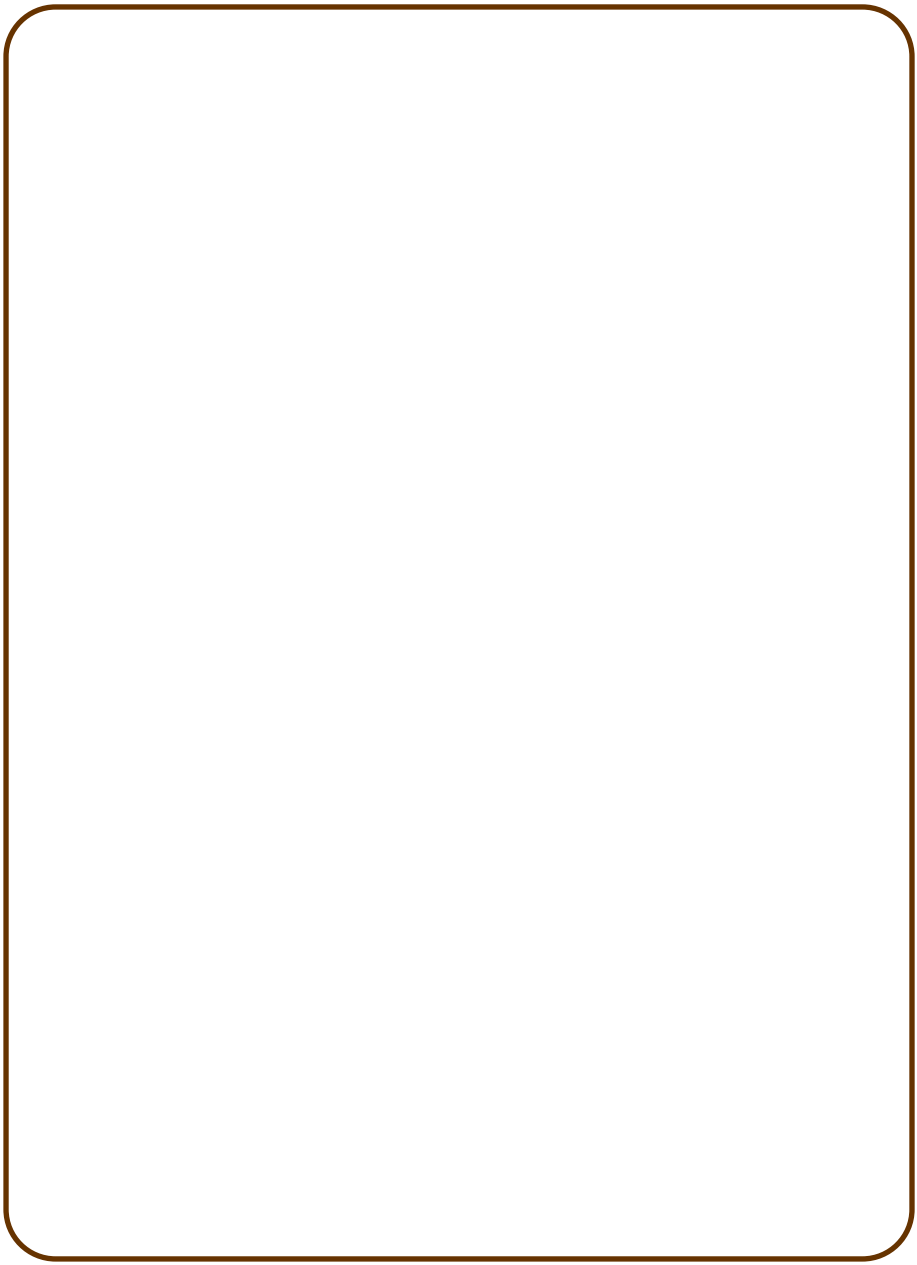

SIGMOD/PODS 2012 ON SOCIAL MEDIA

Follow us on Twitter at http://twitter.com/acmsigmod2012

Please use the hashtags **#sigmod12** and/or **#pods12** to discuss in general about the conference. We have assigned specific hashtags for each paper, and encourage you to use them. There is a simple, session-based algorithm to derive a paper's hashtag, which, in conjunction with #sigmod12 and #pods12 will uniquely identify a paper:

SIGMOD Research Track Papers	**#sigmod12 #RxPy** where x is the session number (1–16) and y is the paper number within the session
SIGMOD Industrial Track Papers	**#sigmod12 #IxPy** where x is the session number (1–6) and y is the paper number within the session
SIGMOD Demos	**#sigmod12 #demoXY** where X is the session letter (A–C), and Y is the demo number within the session
SIGMOD Tutorials	**#sigmod #tutorialX** where X is the tutorial number (1–6)
SIGMOD Keynotes	**#sigmod12 #keynoteX** where X is the keynote number (1–2)
SIGMOD Plenary Session: Perspectives on Big Data	**#sigmod12 #bigdata**
SIGMOD New Researcher Symposium	**#sigmod12 #nrs**
SIGMOD Undergraduate Research Poster Competition	**#sigmod12 #urpc**
SIGMOD Information Session on Careers in Industry	**#sigmod12 #career**
SIGMOD Award Talks Session	**#sigmod12 #award**
SIGMOD Programming Contest	**#sigmod12 #pc**
Research Poster Plenaries	**#sigmod12 #pods12 #rp1** (combined research plenary) **#sigmod12 #rp2** (Poster Session for Workshop Papers) **#sigmod12 #rp3** (SIGMOD Plenary Poster Session)
PODS Papers and Tutorials	**#pods12 #RxPy** where x is the session number (1–10) and y is the number of the presentation within the session
PODS Keynote	**#pods12 #keynote**