



acm

ACM SIGMOD/PODS 2015

CONFERENCE PROGRAM

May 31 - June 4, 2015

Melbourne, Victoria, Australia

<http://www.sigmod2015.org/>





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

SIGMOD/PODS 2015

**May 31 - June 4
Melbourne, Victoria, Australia
<http://www.sigmod2015.org/>**

Conference Program

Copyright © 2015

Editors: Timos Sellis (RMIT University, Australia), Tova Milo (Tel Aviv University, Israel), Susan B. Davidson (University of Pennsylvania, USA), Zack Ives (University of Pennsylvania, USA), Diego Calvanese (Free University of Bozen-Bolzano, Italy), Huiping Cao (New Mexico State University), Xiang Lian (University of Texas - Pan American), Marco Montali (Free University of Bozen-Bolzano, Italy)

Template: Yannis Velegrakis, Toronto 2004, Available for download at: <http://db.disi.unitn.eu/pages/programBooklet>
Permission to use the template is granted, with the kind request to include the current notice in the booklet.

TABLE OF CONTENTS

[Welcome Message From the SIGMOD Chairs](#)

[Welcome Message From the PODS General and Program Chairs](#)

[SIGMOD Conference Organization](#)

[PODS Conference Organization](#)

[Sponsors, Supporters, and Exhibitors](#)

[Conference Venue](#)

[Program at a Glance](#)

[Session Contents](#)

[Awards](#)

[Keynotes](#)

[Tutorials](#)

[SIGMOD New Researcher Symposium](#)

[SIGMOD Research Paper Abstracts](#)

[SIGMOD Industrial Paper Abstracts](#)

[SIGMOD Demonstration Abstracts](#)

[PODS Paper Abstracts](#)

[SIGMOD Travel Awards](#)

[Ph.D. Workshop Participants](#)

[Co-Located Workshops](#)

[Local Information](#)

[SIGMOD/PODS 2015 on Social Media](#)

WELCOME MESSAGE FROM THE SIGMOD CHAIRS

Welcome to SIGMOD 2015 – officially, the 2015 ACM SIGMOD International Conference on the Management of Data! This year’s conference is being held in the beautiful cultural capital of Australia, Melbourne. During the Gold Rush period of the 19th Century, Melbourne was the richest city in the world, and as a result it is filled with many unique neighborhoods and distinctive buildings. In addition to wonderful neighborhoods to explore, the city has great museums and other cultural attractions, as well as a fine multi-cultural atmosphere. For those who would like to explore the outdoors, popular highlights are the Phillip Island Nature Park (90 minutes away), which features wild penguins who return in a parade each day at sunset, and the Great Ocean Road, one of the world’s most scenic coastal drives, including the famous towering 12 Apostles.

SIGMOD 2015’s exciting technical program reflects not only traditional topics, but the database community’s role in broader data science and data analytics. The keynote from Laura Haas, “The Power Behind the Throne: Information Integration in the Age of Data-Driven Discovery” highlights the role of database and data integration techniques in the growing field of data science. Jignesh Patel’s talk, “From Data to Insights @ Bare Metal Speed,” explains how hardware and software need to be co-evolved to support the needs of scalable data analytics. Jennifer Widom, winner of the 2015 ACM-W Athena Lecturer Award for fundamental contributions to computer science, will give her award talk, “Three Favorite Results,” on Tuesday. Christopher Ré will lead a panel on “Machine Learning and Databases: The Sound of Things to Come or a Cacophony of Hype?,” with participants Divyakant Agrawal, Magdalena Balazinska, Michael Cafarella, Michael Jordan, Tim Kraska, and Raghu Ramakrishnan. Of course, there are also 106 research paper presentations, 4 tutorials, 30 demonstrations, and 18 industrial papers. Papers will be presented both as talks during the research sessions, and as part of plenary Poster Sessions.

SIGMOD 2015 is preceded by the PhD Workshop, as well as workshops on leading-edge topics like data analytics (DanaC), databases and the Web (WebDB), exploratory search (ExploreDB), managing and mining spatial data (GeoRich), and graph data (GRADES); the New Researcher Symposium will take place on Wednesday. The banquet will be held in a Melbourne landmark, the Town Hall.

We are grateful to the many researchers who have shaped the conference program. Specific thanks go to the keynotes chair, Renée Miller; the panel chairs, Michael Franklin and Tova Milo; the tutorial chair, Wenfei Fan; the industrial program chairs, Sihem Amer-Yahia and Christopher Olston, as well as the industrial program committee members; the demonstrations chairs, Vanessa Braganholo and Juliana Freire, as well as the demonstration program committee members; the workshops chairs, Anastasia Ailamaki and Yi Chen; the undergraduate research chairs, Anisoara Nica and Anastasios Kementsietsidis; and the New Researcher Symposium chairs, Julia Stoyanovich and Alexandria Meliou. We would also like to express our immense gratitude to the members of the SIGMOD 2015 program committee and especially the “Group Leaders” who helped oversee the review process: Alin Deutsch, Yanlei Diao, Jens Dittrich, Luna Dong, Hank Korth, Alexandros Labrinidis, Feifei Li, Dan Olteanu, Jignesh Patel, Sunita Sarawagi, Pierre Senellart, Wang-Chiew Tan, Jun Yang, and Xiaofang Zhou.

As in recent years, we had two submission deadlines for SIGMOD this year, one in August and one in November. The review process was journal-style, with multiple rounds of reviews coordinated by the Group Leaders. We accepted 34 of 137 papers from the first deadline and 72 of 278 from the second deadline. The total acceptance rate was about 25.5%, and we believe that the revision process has improved the quality of the technical program.

We thank the members of the SIGMOD 2015 organizing committee, as well as the SIGMOD Executive Committee, for their invaluable help and guidance throughout the process leading up to the conference. The team includes the finance chair Xiuzhen Jenny Zhang, the publicity and social media chairs Egemen Tanin and Jagan Sankaranarayanan, the sponsorship chairs Nick Koudas and Mohamed Mokbel, the local arrangements chair Yanchun Zhang, the registration chair Arkady Zaslavsky, the demonstration and workshop local arrangements chair Benjamin Rubinstein, the web/information chair Kai (Kevin) Zheng, the travel award committee chairs Alexandra Meliou and Eugene Wu, the programming contest chair Thomas Neumann, and the wonderful team of ICE Australia who supported the whole organization. We also extend our appreciation to the student volunteers – the success of this conference is impossible without their

hard work and dedication. Finally, we benefited from a great deal of advice and help from Tamer Ozsu, as a former Program Chair, and K. Selçuk Candan, as both a liaison to the SIGMOD Executive and as a former SIGMOD General Chair. 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 thank the many SIGMOD 2015 supporters, whose contributions helped to keep the conference registration fees manageable, particularly for student attendees. Sponsors include: Platinum Level – Microsoft and Oracle; Gold Level – Facebook, Google, IBM Research, SAP, Tableau, and Twitter; and Silver Level – Intel and RMIT University. NEC has also supported SIGMOD 2015. In addition, Microsoft, Google, the National Science Foundation and the SIGMOD Executive Committee have generously funded the SIGMOD 2015 student travel award program. Google's donation enabled us to offer 2 travel awards, Microsoft's donation provided 18 travel awards, and the NSF grant covered another 18 travel awards. Finally, the SIGMOD EC provided 27 awards and covered all 65 free registrations for the student travel program awardees.

Welcome to Melbourne, and we hope you enjoy the conference!

Timos Sellis
General Chair

Susan Davidson and Zachary Ives
Program Co-Chairs

Huiping Cao and Xiang Lian
Proceedings Co-Chairs

WELCOME MESSAGE FROM PODS GENERAL AND PROGRAM CHAIRS

It is our great pleasure to welcome you to the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS 2015), held in Melbourne, Victoria, Australia, on May 31 – June 4, 2015, in conjunction with the 2015 ACM SIGMOD International Conference on Management of Data. Since the first edition of the symposium in 1982, the PODS papers are distinguished by a rigorous approach to widely diverse problems in data management, 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 PODS Symposia study data management challenges in a variety of application contexts, including more recently probabilistic data, streaming data, graph data, information retrieval, ontology and semantic web, and data-driven processes and systems. PODS has a tradition of being the premier international conference on the theoretical and foundational aspects of data management, and the interested reader is referred to the PODS web pages at <http://www.sigmod.org/the-pods-pages/> for information on the history of this conference series.

This year's symposium continues this tradition, but in addition the PODS Executive Committee decided to broaden the scope of PODS, and to explicitly invite for submission papers providing original, substantial contributions in one or more of the following categories:

- a) deep theoretical exploration of topical areas central to data management;
- b) new formal frameworks that aim at providing the basis for deeper theoretical investigation of important emerging issues in data management; and
- c) validation of theoretical approaches from the lens of practical applicability in data management.

The proceedings of PODS 2015 include an abstract for the keynote address by Michael I. Johnson (University of California, Berkeley), papers based on two invited tutorials by Todd J. Green (LogicBlox, USA) and Graham Cormode (University of Warwick, UK), and 25 contributions that were selected by the Program Committee for presentation at the symposium.

This year, PODS experimented for the first time with two submission cycles, where the first cycle allowed also for papers to be revised and resubmitted. For the first cycle, 29 papers were submitted, 4 of which were directly selected for inclusion in the proceedings, and 7 were invited for a resubmission after a revision. The quality of most of the revised papers increased substantially with respect to the first submission, and 6 of those in the end were selected for the proceedings. For the second cycle, 51 papers were submitted, 15 of which were selected, resulting in 25 papers selected overall from a total number of 80 submissions. Most of the 25 accepted papers are extended abstracts. While all submissions have been reviewed by at least four Program Committee members, they have not been formally refereed. It is expected that much of the research described in these papers will be published in a more polished and detailed form in scientific journals.

With respect to the three categories mentioned above, of the 80 submissions (resp., 25 accepted papers), 47 (resp., 19) were classified by the authors in category (a), 28 (resp., 5) in category (b), and only 6 (resp., 3) in category (c). The categories are non-exclusive, and classification was not mandatory; indeed, several papers were classified in more than one category, and for 13 (resp., 3) submissions, no category was specified.

An important task for the Program Committee has been the selection of the PODS 2015 Best Paper Award. The committee selected the paper *“Parallel-Correctness and Transferability for Conjunctive Queries”* by Tom J. Ameloot, Gaetano Geck, Bas Ketsman, Frank Neven, and Thomas Schwentick. On behalf of the committee, we would like to extend our sincere congratulations to the authors.

Since 2008, PODS assigns the ACM PODS Alberto O. Mendelzon Test-of-Time Award to a paper or a small number of papers published in the PODS proceedings ten years prior that had the most impact over the intervening decade. This year's committee, consisting of Dan Suciu (chair), Foto Afrati, and Frank Neven, selected the following two papers. Our warmest congratulations to their authors!

“XPath Satisfiability in the Presence of DTDs”, by Michael Benedikt, Wenfei Fan, and Floris Geerts
“Views and Queries: Determinacy and Rewriting”, by Luc Segoufin and Victor Vianu.

We thank all authors who submitted papers to the symposium. We would also like to thank all members of the Program Committee and all external referees for the enormous amount of work they have done. The two submission cycles, with the revision round for the first cycle, placed an extra burden on all of them. The program committee did not meet in person, but carried out extensive discussions during the electronic program committee meeting, relying on the EasyChair System.

We thank Maurizio Lenzerini and Richard Hull, the former two PODS General Chairs, for their unfailing support. Our thanks go also to the members of the PODS Executive Committee, who advised in the selection of the Program Committee, and who gave important advise on policy issues in all phases of the conference organization. We are particularly grateful to Martin Grohe, the previous PODS PC chair, for his help and advise. Special thanks go to Marco Montali, the Proceedings and Publicity Chair of PODS 2015, and to Wim Martens, for maintaining the PODS Web pages. We are also grateful to the SIGMOD 2015 General Chair, Timos Sellis, for his excellent coordination of the overall SIGMOD/PODS 2015 conference, and his collaboration and support in all the issues requiring coordination between SIGMOD and PODS. Finally, we thank all our sponsors, in particular the ACM Special Interest Groups on Management of Data, on Algorithms and Computation Theory, and on Artificial Intelligence, for their invaluable support.

Tova Milo

PODS 2015 General Chair

Diego Calvanese

PODS 2015 Program Chair

SIGMOD CONFERENCE ORGANIZATION

General Chair: Timos Sellis (RMIT University, Australia)

Program Chairs: Susan B. Davidson (University of Pennsylvania, USA)
Zack Ives (University of Pennsylvania, USA)

Proceedings Chairs: Huiping Cao (New Mexico State University, USA)
Xiang Lian (University of Texas - Pan American, USA)

Tutorial Chair: Wenfei Fan (University of Edinburgh, UK)

Keynotes Chair: Renée Miller (University of Toronto, Canada)

Panels Chairs: Michael Franklin (University of California, Berkeley, USA)
Tova Milo (Tel Aviv University, Israel)

Industrial Program Chairs: Sihem Amer-Yahia (CNRS, Grenoble, France)
Christopher Olston (Google, USA)

Demonstration Chairs: Vanessa Braganholo (Universidade Federal Fluminense, Brazil)
Juliana Freire (New York University, USA)

Workshops Chairs: Anastasia Ailamaki (Ecole Polytechnique Federale de Lausanne, Switzerland)
Yi Chen (New Jersey Institute of Technology, USA)

Undergraduate Research Program Chairs: Anisoara Nica (SAP)
Anastasios Kementsietsidis (Google)

New Researcher Symposium Chairs: Julia Stoyanovich (Drexel University, USA)
Alexandra Meliou (University of Massachusetts Amherst, USA)

Finance Chair: Xiuzhen Jenny Zhang (RMIT University, Australia)

Publicity and Social Media Chairs: Egemen Tanin (University of Melbourne, Australia)
Jagan Sankaranarayanan (NEC Labs America, USA)

Sponsorship Chairs: Nick Koudas (University of Toronto, Canada)
Mohamed Mokbel (University of Minnesota, USA)

Programming Contest Chair: Thomas Neumann (Technische Universität München, Germany)

Local Arrangements Chair: Yanchun Zhang (Victoria University, Australia)

Registration Chair: Arkady Zaslavsky (CSIRO, Australia)

Demonstration, Workshop, and Local Arrangements Chair: Benjamin Rubinstein (University of Melbourne, Australia)

Travel Award Selection Committee Chairs: Alexandra Meliou (University of Massachusetts Amherst, USA)
Eugene Wu (Massachusetts Institute of Technology, USA)

Web/Information Chair: Kai (Kevin) Zheng (The University of Queensland, Australia)

Program Committee Group Leaders: Alin Deutsch (University of California, San Diego, USA)
Yanlei Diao (University of Massachusetts, Amherst, USA)
Jens Dittrich (Saarland University, Germany)
Luna Dong (Google, USA)
Hank Korth (Lehigh University, USA)
Alex Labrinidis (University of Pittsburgh, USA)
Feifei Li (University of Utah, USA)
Dan Olteanu (University of Oxford, UK)
Jignesh Patel (University of Wisconsin, USA)
Sunita Sarawagi (Indian Institute of Technology, Bombay, India)
Pierre Senellart (Telecom ParisTech, France)
Wang-Chiew Tan (University of California, Santa Cruz, USA)
Jun Yang (Duke University, USA)
Xiaofang Zhou (University of Queensland, Australia)

Program Committee: Yanif Ahmad (John Hopkins University, USA)
Reza Akbarinia (INRIA Sophia Antipolis, France)
Magda Balazinska (University of Washington, USA)
Zhuwei Bao (Facebook, USA)
Pablo Barcelo (Universidad de Chile, Chile)
Srikanta Bedathur (Indraprastha Institute of Information Technology, Delhi, India)
Spyros Blanas (Ohio State University, USA)
Philip Bohannon (Facebook, USA)
Stephane Bressan (National University of Singapore, Singapore)
Michael Cafarella (University of Michigan, USA)
K. Selçuk Candan (Arizona State University, USA)
Badrish Chandramouli (Microsoft Research, USA)
Shimin Chen (Chinese Academy of Science, China)
Lei Chen (The Hong Kong University of Science & Technology, China)
James Cheng (The Chinese University of Hong Kong, China)
Alvin Cheung (Massachusetts Institute of Technology, USA)
Fei Chiang (McMaster University, Canada)
Rada Chirkova (North Carolina State University, USA)
John Cieslewicz (Google, USA)
Sarah Cohen-Boulakia (LRI Université Paris Sud, France)
Dario Colazzo (Université Paris-Dauphine, France)
Tyson Condie (University of California, Los Angeles, USA)

Program Committee (continued):

Bin Cui (Peking University, China)
 Antonios Deligiannakis (Technical University of Crete, Greece)
 Amol Deshpande (University of Maryland, USA)
 Daniel Deutch (Tel Aviv University, Israel)
 Eduard Dragut (Temple University, USA)
 Alan Fekete (University of Sydney, Australia)
 Avriilia Floratou (IBM Almaden, USA)
 Tim Furche (University of Oxford, UK)
 Minos Garofalakis (Technical University of Crete, Greece)
 Wolfgang Gatterbauer (Carnegie Mellon University, USA)
 Tingjian Ge (University of Massachusetts, Lowell, USA)
 Boris Glavic (Illinois Institute of Technology, Chicago, USA)
 TJ Green (LogicBlox, USA)
 Torsten Grust (University of Tübingen, Germany)
 Carmem Hara (Universidade Federal do Parana, Brazil)
 Michael Hay (Colgate University, USA)
 Melanie Herschel (University Paris-Sud 11, France)
 Vagelis Hristidis (University of California, Irvine, USA)
 Panos Ipeirotis (New York University, USA)
 Chris Jermaine (Rice University, USA)
 Alekh Jindal (Massachusetts Institute of Technology, USA)
 Panagiotis Karras (Skoltech, Russia)
 Gjergji Kasneci (Hasso-Plattner-Institute, Germany)
 Tasos Kementsietsidis (IBM Research, USA)
 Oliver Kennedy (University at Buffalo, USA)
 Martin Kersten (Centrum Wiskunde & Informatica, Netherlands)
 Daniel Kifer (Pennsylvania State University, USA)
 George Kollios (Boston University, USA)
 Tim Kraska (Brown University, USA)
 Wolfgang Lehner (Technische Universität Dresden, Germany)
 Chengkai Li (University of Texas, Arlington, USA)
 Guoliang Li (Tsinghua University, China)
 Lipyeeow Lim (University of Hawaii at Manoa, USA)
 Xuemin Lin (University of New South Wales, Australia)
 Chengfei Liu (Swinburne University of Technology, Australia)
 Jiaheng Lu (Renmin University of China, China)
 Qiong Luo (The Hong Kong University of Science and Technology, China)
 Ashwin Machanavajjhala (Duke University, USA)
 Nikos Mamoulis (University of Hong Kong, China)
 Silviu Maniu (University of Hong Kong, China)
 Amelie Marian (Rutgers University, USA)
 Davide Martinenghi (Politecnico di Milano, Italy)
 Sharad Mehrotra (University of California, Irvine, USA)
 Alexandra Meliou (University of Massachusetts, USA)
 Xiaofeng Meng (Renmin University of China, China)
 Ahmed Metwally (Google, USA)
 Sebastian Michel (Saarland University, Germany)
 Arnab Nandi (Ohio State University, USA)
 Thomas Neumann (Technische Universität München, Germany)
 Fatma Özcan (IBM research, USA)
 Olga Papaemmanouil (Brandeis University, USA)
 Stelios Pappas (Microsoft Research, USA)
 Aditya Parameswaran (University of Illinois at Urbana-Champaign, USA)
 Srinivasan Parasarathy (Ohio State University, USA)
 Andy Pavlo (Carnegie Mellon University, USA)
 Rachel Pottinger (University of British Columbia, Canada)
 Nicoleta Preda (University of Versailles, France)
 Li Qian (Facebook)

Program Committee (continued):

Jorge Quiané-Ruiz (Qatar Computing Research Institute, Qatar)
Rajeev Rastogi (Amazon, USA)
Chris Ré (Stanford University, USA)
Juan Reutter (Pontificia Universidad Católica de Chile, Chile)
Sudeepa Roy (University of Washington, USA)
Senjuti B Roy (University of Washington Tacoma, USA)
Elke A Rundensteiner (Worcester Polytechnic Institute, USA)
Florin Rusu (University of California, Merced, USA)
Cyrus Shahabi (University of Southern California, USA)
Mohamed A Sharaf (University of Queensland, Australia)
Dennis Shasha (New York University, USA)
Michael Spear (Lehigh University, USA)
Julia Stoyanovich (Drexel University, USA)
Kian-Lee Tan (National University of Singapore, Singapore)
Nan Tang (Qatar Computing Research Institute, Qatar)
Yufei Tao (The Chinese University of Hong Kong, China)
Nesime Tatbul (Massachusetts Institute of Technology, USA)
Evimaria Terzi (Boston University, USA)
Martin Theobald (University of Antwerp, Belgium)
Dimitrios Theodoratos (New Jersey Institute of Technology, USA)
Dimitrios Tsoumakos (National Technical University of Athens, Greece)
Vasilis Vassalos (Athens University of Economics and Business, Greece)
Marcos Vaz Salles (University of Copenhagen - DIKU, Denmark)
Stratis Viglas (University of Edinburgh, UK)
Haixun Wang (Google Inc., USA)
Steven E Whang (Google Inc., USA)
Sai Wu (Zhejiang University, China)
Yinghui Wu (University of California, Santa Barbara, USA)
Melanie Y Wu (Indiana University, USA)
Kuixiao Xiao (Nanyang Technological University, Singapore)
Jianliang Xu (Hong Kong Baptist University, China)
Ke Yi (The Hong Kong University of Science & Technology, China)
Cong Yu (Google, USA)
Wenyuan Yu (Facebook, USA)
Jeffrey Xu Yu (The Chinese University of Hong Kong, China)
Carlo Zaniolo (University of California, Los Angeles, USA)
Demetris Zeinalipour (University Cyprus, Cyprus)
Rui Zhang (University Melbourne, Australia)
Baihua Zheng (Singapore Management University, Singapore)
Kai Zheng (The University of Queensland, Australia)

**Tutorial Program Committee
Members:**

Philip Bohannon (Facebook, USA)
Floris Geerts (University of Antwerp, Belgium)
Felix Naumann (Hasso Plattner Institut, Germany)
Meral Ozsoyoglu (Case Western Reserve University, USA)
Tamer Özsu (University of Waterloo, Canada)
Jignesh Patel (University of Wisconsin, USA)
Divesh Srivastava (AT&T Labs, USA)
Haixun Wang (Google, USA)

**Demonstration Program
Committee Members:**

Amr El Abbadi (University of California Santa Barbara, USA)
Serge Abiteboul (INRIA, France)
Medha Atre (University of Pennsylvania, USA)
Luciano Barbosa (IBM Research, Brazil)
Yi Chen (New Jersey Institute of Technology, USA)
Sara Cohen (The Hebrew University of Jerusalem, Israel)
Eli Cortez (Microsoft Research, USA)

Lois Delcambre (Portland State, USA)
 Boris Glavic (IIT Chicago, USA)
 Todd Green (LogicBlox, USA)
 Carmem Hara (Universidade Federal do Paraná, Brazil)
 Thomas Heinis (Imperial College, UK)
 Stratos Idreos (Harvard, USA)
 Dmitri V. Kalashnikov (University of California Irvine, USA)
 Verena Kantere (University of Geneva, Switzerland)
 Yannis Katsis (University of California San Diego, USA)
 David Koop (UMass-Dartmouth, USA)
 Flip Korn (Google Research)
 Georgia Koutrika (HP Labs, USA)
 Chao Li (Google Inc., USA)
 Ashwin Machanavajjhala (Duke University, USA)
 Ioana Manolescu (INRIA, France)
 Yannis Manolopoulos (Aristotle University of Thessaloniki, Greece)
 Paolo Meriardo (Universita Roma Tre, Italy)
 Mirella Moro (Universidade Federal de Minas Gerais, Brazil)
 Mario Nascimento (University of Alberta, Canada)
 Eduardo Ogasawara (CEFET/RJ, Brazil)
 Daniel Oliveira (Universidade Federal Fluminense, Brazil)
 Andy Pavlo (Carnegie Mellon University, USA)
 Lucian Popa (IBM Research, USA)
 Maya Ramanath (IIT Delhi, India)
 Sudeepa Roy (University of Washington, USA)
 Yasin N. Silva (Arizona State University, USA)
 Emad Soroush (GraphLab Inc, USA)
 Valéria Times (Universidade Federal de Pernambuco, Brazil)
 Ricardo Torres (UNICAMP, Brazil)
 Pinar Tozun (EPFL, Switzerland)
 Huy Vo (New York University, USA)
 Xiaochun Yang (Northeastern University, China)
 Wenchao Zhou (Georgetown University, USA)

**Industrial Program Committee
 Members:**

Jose Blakeley (Microsoft)
 Paul Brown (Paradigm 4)
 Ed H. Chi (Google, USA)
 Carlo Curino (Microsoft)
 Shel Finkelstein (University of California Santa Cruz)
 Stratos Idreos (Harvard University, USA)
 Mohsen Jamali (WalmartLabs)
 Grigoris Karvounarakis (LogicBlox, Greece)
 Donald Kossmann (ETH Zurich)
 Benjamin Reed (Facebook)
 Jean-Michel Renders (Xerox Labs, France)
 Philippe Rigaux (Internet Memory, France)
 Mehul Shah (Amiatio)
 Eric Simon (SAP, France)
 Theo Vassilakis (Metanautix)
 Todd Walter (Teradata)
 Ken Yocum (Illumina)

**Additional Reviewers (In
 alphabetical order by Last
 Name):**

Afsin Akdogan	Gerasimos Marketos
Cem Aksoy	Thomas Marshall
Peter Alvaro	Julian McAuley
Panayiotis Andreou	Rui Meng
Vaibhav Arora	Xianrui Meng
Joy Arulraj	Umar Farooq Minhas

Mohammad Asghari	Prasenjit Mitra
Ramoza Ashan	Jayanta Mondal
Andreas Behrend	Tobias Müller
Souvik Bhattacharjee	Faisal Nawab
Panagiotis Bouros	Apostolos Papadopoulos
Matteo Brucato	Nikolaos Papailiou
Francesca Bugiotti	Huy Pham
Damian Bursztyn	Yann Ponty
Lei Cao	Olga Poppe
Bogdan Cautis	Mohiuddin Abdul Qader
Aniket Chakrabarty	Li Qian
Lijun Chang	Chengjie Qin
Georgios Chatzimilioudis	Lu Qin
Amit Chavan	Teng Qin
Muhammad Aamir Cheema	Abdul Quamar
Rui Chen	Nisarg Raval
Yan Chen	Medhabi Ray
Shiwen Cheng	Alexandra Roatis
Yu Cheng	Yiye Ruan
Shumo Chu	Shouq Sadah
Eleonora Ciceri	Cetin Sahin
Andrew Crotty	Zechao Sang
Mahashweta Das	Carlo Sartiani
Ananya Dass	Rebeca Schroeder
Dingxiong Deng	Vivek Shah
Dong Deng	Moloud Shahbazi
Aggeliki Dimitriou	Anil Shanbhag
Katerina Doka	Yingxia Shao
S.M. Faisal	Jieming Shi
Liyue Fan	Juwei Shi
Ioannis Flouris	Yanchen Shi
David Fuhry	Leferis Sidirourgos
Alexander Galakatos	Afroza Sultana
Yihan Gao	Youze Tang
Theodore Georgiou	Yu Tang
Nikos Giatrakos	Hien To
Anastasios Gounaris	Yongxin Tong
Eric Gribkoff	Tassos Venetis
Sairam Gurajada	Domagoj Vrgoc
Samuel Haney	Chi Wang
Sona Hasani	Jingjing Wang
Naeemul Hassan	Sibo Wang
Xi He	Wei Wang
Zhenying He	Xiang Wang
Haibo Hu	Xiaolan Wang
Huiqi Hu	Xiaoyang Wang
Silu Huang	Zhengkui Wang
Weihuang Huang	Mingrui Wei
Xin Huang	Zhewei Wei
Yuzheng Huang	Zeyi Wen
Zengfeng Huang	Matthew Wiley
Sergey Ivanov	Huanhuan Wu
Srini Iyer	Xiaoying Wu
Ayush Jain	Yingjun Wu
Nandish Jayaram	Qin Xiao
Lilong Jiang	Jie Xu
Niranjan Kamat	Da Yan
Zoi Kaoudi	Fan Yang
Konstantinos Karanasos	Shiyu Yang
Ashwin Kayyoor	Erfan Zamanian

Udayan Khurana
Albert Kim
Georgios Kollias
Ios Kotsogiannis
Christos Laoudias
Chuan Lei
Jianxin Li
Jinfeng Li
Yafei Li
Yaguang Li
Xiang Lian
Xika Lin
Miguel Liroz Gistau
Yi Lu
Ying Lu

Stamatis Zampetakis
Gensheng Zhang
Jun Zhang
Xiaofei Zhang
Yang Zhang
Ying Zhang
Rui Zhou
Haohan Zhu
Linhong Zhu
Honglei Zhuang
Zhongfang Zhuang
George Zois
Zainab Zolaktaf
Andreas Zuefle

PODS CONFERENCE ORGANIZATION

General Chair: Tova Milo (Tel Aviv University, Israel)

Program Chair: Diego Calvanese (Free University of Bozen-Bolzano, Italy)

Proceedings and Publicity Chair: Marco Montali (Free University of Bozen-Bolzano, Italy)

Program Committee Members: Peyman Afshani (Aarhus University, Denmark)
Edith Cohen (Microsoft Research, USA)
Minos Garofalakis (Technical University of Crete, Greece)
Lukasz Golab (University of Waterloo, Canada)
Gianluigi Greco (University of Calabria, Italy)
Benny Kimelfeld (LogicBlox, USA)
Leonid Libkin (University of Edinburgh, UK)
Carsten Lutz (University of Bremen, Germany)
Ioana Manolescu (INRIA Saclay Ile-de-France, France)
Yakov Nekrich (University of Waterloo, Canada)
Dan Olteanu (University of Oxford, UK)
Jeff M. Phillips (University of Utah, USA)
Reinhard Pichler (Technical University Vienna, Austria)
Riccardo Rosati (Sapienza University of Rome, Italy)
Mantas Simkus (Technical University Vienna, Austria)
Cristina Sirangelo (ENS Cachan, France)
Jianwen Su (University of California Santa Barbara, USA)
Yannis Theodoridis (University of Pyraeus, Greece)
Stijn Vansummeren (University Libre De Bruxelles, Belgium)
Rossano Venturini (University of Pisa, Italy)
Limsoon Wong (National University of Singapore, Singapore)

Additional reviewers:

Amir Abboud	Andrew McGregor
Mario Alviano	Gerome Miklau
Mohammad Amiri	Gonzalo Navarro
Alexandr Andoni	Frank Neven
Myrto Arapinis	Magdalena Ortiz
Arvind Arasu	Rasmus Pagh
David Balduzzi	Nikos Pelekis
Luca Becchetti	Jorge Perez
Andrea Cali	Andreas Pieris
Karthekeyan	Christopher Ré
Chandrasekaran	Stefan Rmmele
Giuseppe De Giacomo	Raef Raef
Daniel Deutch	Ilya Razenshteyn
David Dominguez-Sal	Dana Ron
Christos Doulkeridis	Emanuel Sallinger
Cynthia Dwork	Swaminathan Sankararaman
David Eppstein	Vadim Savenkov
Andrea Farruggia	Francesco Scarcello
Rainer Gemulla	Sebastian Skritek
Mina Ghashami	Christian Sohler
André Hernich	Uri Stemmer
Xiaocheng Hu	Xiaoming Sun
Zengfeng Huang	Yutian Sun
T. S. Jayram	Yan Tang
Hossein Jowhari	Yufei Tao
Shivaram Kalyanakrishnan	Balder ten Cate
Michael Kapralov	Abhradeep Guha Thakurta
Eamonn Keogh	Ari Trachtenberg
Michal Kratky	Suresh Venkatasubramanian

Additional reviewers (continued):

Martin Lackner
Domenico Lembo
Maurizio Lenzerini
Ninghui Li
Zhenming Liu
Ashwin Machanavajjhala

Stratis Viglas
David Woodruff
Ke Yi
Haifeng Yu
Qin Zhang

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

Foto Afrati (National Technical University of Athens, Greece)
Frank Neven (Hasselt University, Belgium)
Dan Suciu (chair) (University of Washington, USA)

SPONSORS, SUPPORTERS, AND EXHIBITORS

Sponsor:



Platinum Sponsors:



Gold Sponsors:



Silver Sponsors:



Other Sponsor:



**Platinum
Publisher:**



**Gold
Publisher:**



CONFERENCE VENUE

Melbourne Convention & Exhibition Centre



The 2015 ACM SIGMOD/PODS Conference will be held at the Melbourne Convention and Exhibition Centre. Recognised as Australasia's Leading Meetings and Conference Centre in 2014, 2013 and 2012 by the prestigious World Travel Awards, MCEC proudly offers delegates amazing spaces, leading technology, creative food and wine and staff with the skills and experience to bring it all together. The Convention Centre is just a few minutes walk to the heart of Melbourne and there is plenty to explore in the immediate vicinity.

TRANSPORT OPTIONS

Arriving by taxi

Taxi drop off points are at 1 Convention Centre Place (for Convention Centre) or 2 Clarendon Street, South Wharf (for Exhibition Centre). Taxis can be booked by contacting Silver Top Taxi (131 008) or Yellow Cabs (13 22 27).

Arriving by tram

Melbourne's tram network services the metropolitan inner-city area with 28 lines. Melbourne Convention and Exhibition Centre (MCEC) is serviced by regular services departing every few minutes.

MCEC is stop 124a (Crown Casino / MCEC) for routes 96, 109 and 112 stop opposite the MCEC on Clarendon Street, South Wharf. Tram numbers 48 and 70 stop at the end of Flinders Street. Walk towards the Yarra River and cross Seafarers Bridge.

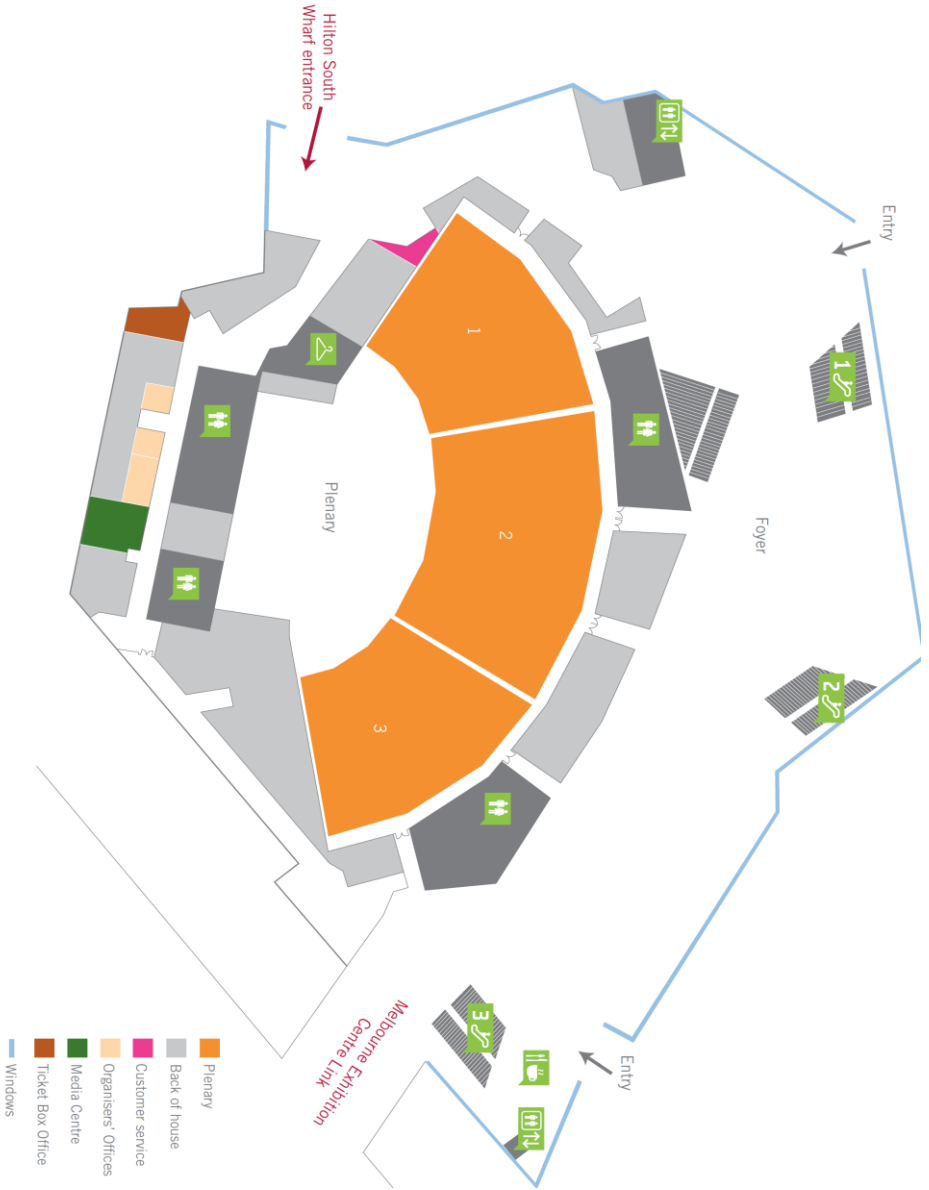
Arriving by train

Melbourne's metropolitan train network operates from the nearby Southern Cross station connecting to all parts of Melbourne. Exit onto Spencer Street and take trams 96, 109 or 112 to stop 124a, which stops opposite the Exhibition Centre entrance of MCEC.

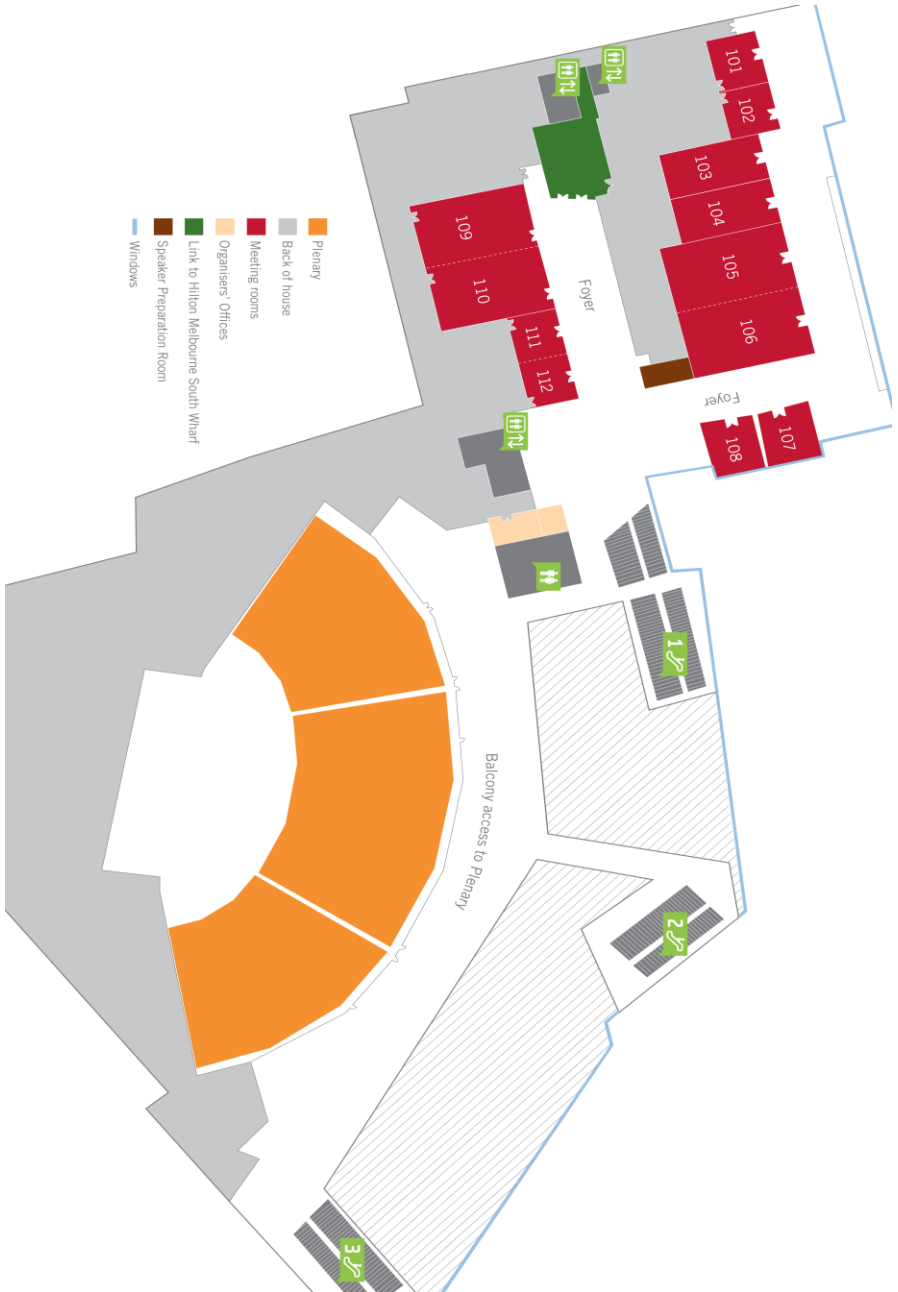
Arriving by bus

Bus route 237 operates from Queen Victoria Market, via Southern Cross Station to Lorimer Street South Wharf from Monday to Friday. Lorimer Street is approximately a five minute walk to MCEC. There is also a coach pick up/drop off point at Bay 1, Convention Centre Place (closest to DFO South Wharf).

CONVENTION CENTRE MAP – GROUND FLOOR



CONVENTION CENTRE MAP – LEVEL 1



PROGRAM AT A GLANCE

SUNDAY

8:30-10:00	GeoRich MR101&102	GRADES MR103	DanaC MR104	WebDB MR105	PhD Workshop MR106	ExploreDB MR107
10:00-10:30	Coffee Break Level 1 Foyer					
10:30-12:10	GeoRich MR101&102	GRADES MR103	DanaC MR104	WebDB MR105	PhD Workshop MR106	ExploreDB MR107
12:10-13:30	Lunch Break (lunch not provided by conference)					
13:30-15:10	GeoRich MR101&102	GRADES MR103	DanaC MR104	WebDB MR105	PhD Workshop MR106	ExploreDB MR107
15:10-15:40	Coffee Break Level 1 Foyer					
15:40-17:20	GeoRich MR101&102	GRADES MR103	DanaC MR104	WebDB MR105	PhD Workshop MR106	ExploreDB MR107
18:00-20:00	PODS Welcome Reception Level 1 Foyer					

MONDAY

07:30-08:15	Morning Tea and Coffee Level 1 Foyer					
8:30-10:00	PODS Keynote 1 MR105 & MR106					
10:00-10:30	Coffee Break Level 1 Foyer					
10:30-12:10	DaMoN Opening and Keynote MR103	PODS Session 1 (ending at 12:00) MR105 & MR106				
12:10-13:30	Lunch and Posters Level 1 Foyer					
13:30-14:30	DaMoN Session 1 MR103	PODS Session 2 MR105 & MR106				
14:30-14:45	Break Level 1 Foyer					
14:45-16:00	DaMoN Session 2 MR103	PODS Session 3 – Tutorial Session 1 MR105 & MR106				
16:00-16:30	Coffee Break Level 1 Foyer					
16:30-18:30	DaMoN Session 3 MR103	PODS Session 4 MR105 & MR106				
18:30-19:30	PODS Business Meeting MR105 & MR106					
19:15-21:15	SIGMOD Welcome Reception + Undergraduate Research Poster Session Polly Woodside Tall Ship reception area					

TUESDAY

07:30-08:15	Morning Tea and Coffee <i>Level 1 Foyer</i>					
08:20-10:00	SIGMOD Keynote Talk 1 <i>Plenary 1</i>					
10:00-10:30	Coffee Break <i>Level 1 Foyer</i>					
10:30-12:10	PODS Session 5 (ending at 12:00) <i>Plenary 1</i>	SIGMOD Research 1 <i>MR106</i>	SIGMOD Research 2 <i>MR101&102</i>	SIGMOD Research 3 <i>MR104</i>	SIGMOD Industry 1 <i>MR103</i>	SIGMOD Tutorial 1.1 <i>MR105</i>
12:10-13:30	Lunch and Posters <i>Level 1 Foyer</i>					
13:30-15:10	PODS Session 6 <i>Plenary 1</i>	Panel <i>MR106</i>	SIGMOD Research 4 <i>MR101&102</i>	SIGMOD Research 5 <i>MR103</i>		SIGMOD Tutorial 1.2 <i>MR105</i>
15:10-15:40	Coffee Break <i>Level 1 Foyer</i>					
15:40-17:20	PODS Session 7 (starting at 15:30) <i>Plenary 1</i>	SIGMOD Research 6 <i>MR101&102</i>	SIGMOD Research 7 <i>MR105</i>	SIGMOD Research 8 <i>MR104</i>	SIGMOD Industry 2 <i>MR103</i>	SIGMOD Programming Contest <i>MR106</i>
17:20-17:30	Break					
17:30-18:15	ACM-W Athena Lecturer Award <i>Plenary 1</i>					
18:15-18:25	Break					
18:25-18:55	SIGMOD Business Meeting <i>Plenary 1</i>			Careers in Industry <i>MR106</i>		
18:55-19:30						

WEDNESDAY

07:30-08:15	Morning Tea and Coffee <i>Level 1 Foyer</i>					
08:20-10:00	SIGMOD Keynote Talk 2 <i>Plenary 1</i>					
10:00-10:30	Coffee Break <i>Level 1 Foyer</i>					
10:30-12:10	PODS Session 8 (ending at 12:00) <i>Plenary 1</i>	SIGMOD Research 9 MR103	SIGMOD Research 10 MR104	SIGMOD Research 11 MR101&102	SIGMOD Demo A MR106	SIGMOD Tutorial 2.1 MR105
12:10-13:20	Lunch and Posters <i>Level 1 Foyer</i>					
13:20-14:50	New Researchers Symposium <i>Plenary 1</i>					
14:50-15:20	Coffee Break <i>Level 1 Foyer</i>					
15:20-17:00	PODS Session 9 (ending at 16:50) <i>Plenary 1</i>	SIGMOD Research 12 MR103	SIGMOD Research 13 MR104		SIGMOD Demo B MR106	SIGMOD Tutorial 2.2 MR105
17:00-17:10	Break					
17:10-18:50	PODS Session 10 (ending at 18:40) <i>Plenary 1</i>	SIGMOD Research 14 MR101&102	SIGMOD Research 15 MR103	SIGMOD Research 16 MR104	SIGMOD Demo C MR106	SIGMOD Industry 3 MR105
18:50-19:15	Break					
19:15-22:00	Conference Banquet <i>Melbourne Town Hall</i>					

THURSDAY

07:30-08:15	Morning Tea and Coffee <i>Level 1 Foyer</i>					
08:30-10:00	SIGMOD Awards Talks <i>Planery 1</i>					
10:00-10:30	Coffee Break <i>Level 1 Foyer</i>					
10:30-12:10	SIGMOD Research 17 <i>MR103</i>	SIGMOD Research 18 <i>MR101 & 102</i>	SIGMOD Research 19 <i>MR104</i>	SIGMOD Industry 4 <i>MR105</i>	SIGMOD Demo B <i>MR106</i>	SIGMOD Tutorial 3.1 <i>Planery 1</i>
12:10-13:30	Lunch and Posters <i>Level 1 Foyer</i>					
13:30-15:10	SIGMOD Research 20 <i>MR103</i>	SIGMOD Research 21 <i>MR101 & 102</i>	SIGMOD Research 22 <i>MR104</i>	SIGMOD Industry 5 <i>MR105</i>	SIGMOD Demo C <i>MR106</i>	SIGMOD Tutorial 3.2 <i>Planery 1</i>
15:10-15:40	Coffee Break <i>Level 1 Foyer</i>					
15:40-17:20	SIGMOD Research 23 <i>MR103</i>		SIGMOD Research 24 <i>MR104</i>		SIGMOD Demo A <i>MR106</i>	SIGMOD Tutorial 4 <i>MR105</i>

SESSION CONTENTS

SUNDAY, 18:00 – 20:00

PODS Welcome Reception

Location: Level 1 Foyer

MONDAY, 8:30 – 10:00

PODS Keynote

Location: MR105&106

Hashtag: #pods15 #keynote

Session Chair: Tova Milo (Tel Aviv University)

Computational Thinking, Inferential Thinking and "Big Data"

Michael I. Jordan, University of California, Berkeley

MONDAY, 10:30 – 12:00

PODS Session 1: Repairs and Ontologies

Location: MR105&106

Hashtag: #pods15 #R1Px

Session Chair: Diego Calvanese (Free University of Bozen-Bolzano)

Dichotomies in the Complexity of Preferred Repairs

Ronald Fagin, IBM Research; Benny Kimelfeld, LogicBlox; Phokion Kolaitis, University of California Santa Cruz

The Data Complexity of Consistent Query Answering for Self-Join-Free Conjunctive Queries Under Primary Key Constraints

Paraschos Koutris, University of Washington; Jef Wijsen, University of Mons

High-Level Why-Not Explanations using Ontologies

Balder Ten Cate, LogicBlox; Cristina Civili, Sapienza University of Rome; Evgeny Sherkhonov, University of Amsterdam; Wang-Chiew Tan, University of California Santa Cruz

MONDAY, 13:30 – 14:30

PODS Session 2: Awards Session

Location: MR105&106

Hashtag: #pods15 #R2Px

Session Chair: Tova Milo (Tel Aviv University)

The ACM PODS Alberto O. Mendelzon Test-of-Time Award:

XPath Satisfiability in the Presence of DTDs

Michael Benedikt, Oxford University; Wenfei Fan, University of Edinburgh; Floris Geerts, University of Antwerp

Views and Queries: Determinacy and Rewriting

Luc Segoufin, INRIA & LSV Cachan; Victor Vianu, University of California San Diego

PODS Best Paper Award:

Parallel-Correctness and Transferability for Conjunctive Queries

Tom J. Ameloot, Hasselt University; Gaetano Geck, TU Dortmund University; Bas Ketsman, Hasselt University; Frank Neven, Hasselt University; Thomas Schwentick, TU Dortmund University

MONDAY, 14:45 – 16:00

PODS Session 3: Tutorial Session 1

Location: MR105&106

Hashtag: #pods15 #R3P1

Session Chair: Reinhard Pichler (Vienna University of Technology)

Invited Tutorial: LogiQL: a Declarative Language for Enterprise Applications

Todd J. Green, LogicBlox

MONDAY, 16:30 – 18:30

PODS Session 4: Existential Rules and TGDs

Location: MR105&106

Hashtag: #pods15 #R4Px

Session Chair: Frank Neven (Hasselt University and Transnational University of Limburg)

Function Symbols in Tuple-Generating Dependencies: Expressive Power and Computability

Georg Gottlob, Oxford University; Reinhard Pichler, Vienna University of Technology; Emanuel Sallinger, Vienna University of Technology

Default Negation for Non-Guarded Existential Rules

Mario Alviano, University of Calabria; Andreas Pieris, Vienna University of Technology

Chase Termination for Guarded Existential Rules

Marco Calautti, University of Calabria; Georg Gottlob, Oxford University; Andreas Pieris, Vienna University of Technology

Recovering Exchanged Data

Gosta Grahne, Concordia University; Ali Moallemi, Concordia University; Adrian Onet, Concordia University

MONDAY, 18:30 – 19:30

PODS Business Meeting

Location: MR105&106

MONDAY, 19:15 – 21:15

SIGMOD Welcome Reception + Undergraduate Research Poster Session

Location: Polly Woodside Tall Ship reception area (<http://www.pollywoodside.com.au/>)

TUESDAY, 8:20 – 10:00

SIGMOD Keynote Talk 1

Location: Plenary 1

Hashtag: #sigmod15 #keynote1

Session Chair: Susan Davidson (University of Pennsylvania)

From Data to Insights @ Bare Metal Speed

Jignesh M. Patel, University of Wisconsin-Madison

TUESDAY, 10:30 – 12:10

Location: Plenary 1

Hashtag: #pods15 #R5Px

Session Chair: Jef Wijsen (University of Mons)

The (Almost) Complete Guide to Tree Pattern Containment

Wojciech Czerwiński, University of Warsaw; Wim Martens, University of Bayreuth; Paweł Parys, University of Warsaw; Marcin Przybytko, University of Warsaw

Efficient Evaluation and Approximation of Well-designed Pattern Trees

Pablo Barceló, University of Chile; Reinhard Pichler, Vienna University of Technology; Sebastian Skritek, Vienna University of Technology

BonXai: Combining the Simplicity of DTD with the Expressiveness of XML Schema

Wim Martens, University of Bayreuth; Frank Neven, Hasselt University; Matthias Niewerth, University of Bayreuth; Thomas Schwentick, TU Dortmund University

SIGMOD Research 1: Cloud: Parallel Execution

Location: MR106

Hashtag: #sigmod15 #R1Px

Session Chair: Yanlei Diao (University of Massachusetts, Amherst)

Distributed Outlier Detection using Compressive Sensing

Ying Yan, Microsoft Research; Jiaxing Zhang, Microsoft Research; Bojun Huang, Microsoft Research; Xuzhan Sun, Peking University; Jiaqi Mu, University of Illinois, Urbana-Champaign; Zheng Zhang, NYU Shanghai; Thomas Moscibroda, Microsoft Research

Locality-aware Partitioning in Parallel Database Systems

Erfan Zamanian, Brown University; Carsten Binnig, Baden-Wuerttemberg Cooperative State University; Abdallah Salama, Baden-Wuerttemberg Cooperative

ByteSlice: Pushing the Envelop of Main Memory Data Processing with a New Storage Layout

Ziqiang Feng, The Hong Kong Polytechnic University; Eric Lo, The Hong Kong Polytechnic University; Ben Kao, The University of Hong Kong; Wenjian Xu, The Hong Kong Polytechnic University

Implicit Parallelism through Deep Language Embedding

Alexander Alexandrov, TU Berlin; Andreas Kunft, TU Berlin; Asterios Katsifodimos, TU Berlin; Felix Schöler, TU Berlin; Lauritz Thamsen, TU Berlin; Odej Kao, TU Berlin; Tobias Herb, TU Berlin; Volker Markl, TU Berlin

From Theory to Practice: Efficient Join Query Evaluation in a Parallel Database System

Shumo Chu, University of Washington; Magdalena Balazinska, University of Washington; Dan Suciu, University of Washington

SIGMOD Research 2: Matrix and Array Computations

Location: MR101&102

Hashtag: #sigmod15 #R2Px

Session Chair: Florin Rusu (University of California, Merced)

sPCA: Scalable Principal Component Analysis for Big Data

Tarek Elgamal, Qatar Computing Research Institute; Maysam Yabandeh, Twitter; Ashraf Aboulnaga, Qatar Computing Research Institute; Walid Fathi, NTG Clarity; Mohamed Hefeeda, Qatar Computing Research Institute

Exploiting matrix dependency for efficient distributed matrix computation

Lele Yu, Peking University; Yingxia Shao, Peking University; Bin Cui, Peking University

LEMP: Fast Retrieval of Large Entries in a Matrix Product

Christina Teflioudi, Max Planck Institute for Computer Science; Rainer Gemulla, University of Mannheim; Olga Mykytiuk, Sulzer GmbH

Skew-Aware Join Optimization for Array Databases

Jennie Duggan, Northwestern University; Olga Papaemmanouil, Brandeis University; Leilani Battle, MIT; Michael Stonebraker, MIT

Resource Elasticity for Large-Scale Machine Learning

Botong Huang, Duke University; Matthias Boehm, IBM Research - Almaden; Yuanyuan Tian, IBM Research - Almaden; Berthold Reinwald, IBM Research - Almaden; Shirish Tatikonda, IBM Research - Almaden; Frederick Reiss, IBM Research - Almaden

SIGMOD Research 3: Security and Access Control

Location: MR104

Hashtag: #sigmod15 #R3Px

Session Chair: Lei Chen (The Hong Kong University of Science and Technology)

SEMROD: Secure and Efficient MapReduce over Hybrid Clouds

Kerim Yasin Oktay, University of California, Irvine; Vaibhav Khadilkar, The University of Texas at Dallas; Murat Kantarcioglu, The University of Texas at Dallas; Sharad Mehrotra, University of California, Irvine

Authenticated Online Data Integration Services

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

ENKI: Access Control for Encrypted Query Processing

Isabelle Hang, University of Bremen; Florian Kerschbaum, SAP AG; Ernesto Damiani, University of Milan

Collaborative Access Control in WebdamLog

Vera Zaychik Moffitt, Drexel University; Julia Stoyanovich, Drexel University; Serge Abiteboul, INRIA Saclay & ENS Cachan; Gerome Miklau, University of Massachusetts

Automatic Enforcement of Data Use Policies with DataLawyer

Prasang Upadhyaya, University of Washington; Magdalena Balazinska, University of Washington; Dan Suciu, University of Washington

SIGMOD Industry 1: Streaming/Real-Time/Active

Location: MR 103

Hashtag: #sigmod15 #I1Px

Session Chair: Tyson Condie (University of California, Los Angeles)

TencentRec: Real-time Stream Recommendation in Practice

Yanxiang Huang, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University; Bin Cui, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University; Wenyu Zhang, Tencent Inc.; Jie Jiang, Tencent Inc.; Ying Xu, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University

Twitter Heron: Stream Processing at Scale

Sanjeev Kulkarni, Twitter; Nikunj Bhagat, Twitter; Masong Fu, Twitter; Vikas Kedigehalli, Twitter; Christopher Kellogg, Twitter; Saitesh Mittal, Twitter; Jignesh Patel, University of Wisconsin; Karthik Ramasamy, Twitter; Siddarth Taneja, Twitter

Analytics in Motion - High Performance Event-Processing AND Real-Time Analytics in the Same Database

Lucas Braun, ETH Zurich; Thomas Etter, ETH Zurich; Georgios Gasparis, ETH Zurich; Martin Kaufmann, ETH Zurich; Donald Kossmann, ETH Zurich; Daniel Widmer, ETH Zurich; Aharon Avitzur, Huawei Technologies; Anthony Iliopoulos, Huawei Technologies; Eliezer Levy, Huawei Technologies; Ning Liang, Huawei Technologies

Why Big Data Industrial Systems Need Rules and What We Can Do About It

Paul Suganthan G.C., University of Wisconsin-Madison; Chong Sun, Uber; Krishna Gayatri K., University of Wisconsin-Madison; Haojun Zhang, University of Wisconsin-Madison; Frank Yang, LinkedIn; Narasimhan Rampalli, @WalmartLabs; Shishir Prasad, @WalmartLabs; Esteban Arcaute, @WalmartLabs; Ganesh Krishnan, @WalmartLabs; Rohit Deep, @WalmartLabs; Vijay Raghavendra, @WalmartLabs; AnHai Doan, University of Wisconsin-Madison

SIGMOD Tutorial 1.1: Overview of Data Exploration Techniques

Location: MR105

Hashtag: #sigmod15 #tutorial1

Presenters: Stratos Idreos, Harvard University; Olga Papaemmanouil, Brandeis University; Surajit Chaudhuri, Microsoft Research

TUESDAY, 12:10–13:30

Lunch and Posters

Location: Level 1 Foyer

TUESDAY, 13:30 – 15:10

PODS Session 6: Tutorial Session 2

Location: Plenary 1

Hashtag: #pods15 #R6Px

Session Chair: Pablo Barceló (Universidad de Chile)

Invited Tutorial: Compact Summaries over Large Datasets

Graham Cormode, University of Warwick

Defining Relations on Graphs: How Hard is it in the Presence of Node Partitions?

M Praveen, Chennai Mathematical Institute; B Srivathsan, Chennai Mathematical Institute

SIGMOD Panel - Machine Learning and Databases: The Sound of Things to Come or a Cacophony of Hype?

Location: MR106

Hashtag: #sigmod15 #panel

Panel Chair: Christopher Ré (Stanford University)

SIGMOD Research 4: Cloud: Fault Tolerance, Reconfiguration

Location: MR101&102

Hashtag: #sigmod15 #R4Px

Session Chair: Alexandros Labrinidis (University of Pittsburgh)

Cost-based Fault-tolerance for Parallel Data Processing

Abdallah Salama, Baden-Wuerttemberg Cooperative State University; Carsten Binnig, Baden-Wuerttemberg Cooperative State University; Tim Kraska, Brown University; Erfan Zamanian, Brown University

Squall: Fine-Grained Live Reconfiguration for Partitioned Main Memory Databases

Aaron Elmore, University of Chicago; Vaibhav Arora, University of California, Santa Barbara; Rebecca Taft, MIT; Andrew Pavlo, Carnegie Mellon University; Divyakant Agrawal, University of California, Santa Barbara & Qatar Computing Research Institute; Amr El Abbadi, University of California, Santa Barbara

Madeus: Database Live Migration Middleware under Heavy Workloads for Cloud Environment

Takeshi Mishima, NTT; Yasuhiro Fujiwara, NTT

Lineage-driven Fault Injection

Peter Alvaro, UC Berkeley; Joshua Rosen, UC Berkeley; Joseph Hellerstein, UC Berkeley

SIGMOD Research 5: Keyword Search and Text

Location: MR103

Hashtag: #sigmod15 #R5Px

Session Chair: K. Selçuk Candan (Arizona State University)

Diversity-Aware Top-K Publish/Subscribe for Text Stream

Lisi Chen, Nanyang Technological University; Gao Cong, Nanyang Technological University

Diverse and Proportional Size-l Object Summaries for Keyword Search

Georgios Fakas, HKUST; Zhi Cai, Beijing University of Technology; Nikos Mamoulis, The University of Hong Kong

Local Filtering: Improving the Performance of Approximate Queries on String Collections

Xiaochun Yang, Northeastern University; Yaoshu Wang, University of New South Wales; Bin Wang, Northeastern University; Wei Wang, University of New South Wales

Exact Top-k Nearest Keyword Search in Large Networks

Minhao Jiang, The Hong Kong University of Science and Technology; Ada Fu, The Chinese University of Hong Kong; Raymond Wong, The Hong Kong University of Science and Technology

Efficient Algorithms for Answering the m-Closest Keywords Query

Tao Guo, Nanyang Technological University; Xin Cao, Queen's University Belfast; Gao Cong, Nanyang Technological University

SIGMOD Tutorial 1.2: Overview of Data Exploration Techniques (cont.)

Location: MR105

Hashtag: #sigmod15 #tutorial1

Presenters: Stratos Idreos, Harvard University; Olga Papaemmanouil, Brandeis University; Surajit Chaudhuri, Microsoft Research

TUESDAY, 15:40 – 17:20

PODS Session 7: Query Processing (starting at 15:30)

Location: Plenary 1

Hashtag: #pods15 #R7Px

Session Chair: Dan Suciu (University of Washington)

Querying Big Data by Accessing Small Data

Wenfei Fan, University of Edinburgh; Floris Geerts, University of Antwerp; Yang Cao, University of Edinburgh; Ting Deng, Beihang University; and Ping Lu, Beihang University

Skyline Queries with Noisy Comparisons

Benoit Groz, Paris-Sud University; Tova Milo, Tel-Aviv University

The Communication Complexity of Distributed Set-Joins with Applications to Matrix Multiplication

Dirk Van Gucht, Indiana University; Ryan Williams, Stanford University; David Woodruff, IBM; Qin Zhang, Indiana University

Joins via Geometric Resolutions: Worst-case and Beyond

Mahmoud Abo Khamis, University at Buffalo; Hung Ngo, University at Buffalo; Christopher Re, Stanford University; Atri Rudra, University at Buffalo

SIGMOD Research 6: Graph Primitives

Location: MR101&102

Hashtag: #sigmod15 #R6Px

Session Chair: Ke Yi (The Hong Kong University of Science and Technology)

Minimum Spanning Trees in Temporal Graphs

SILU HUANG, Chinese University of Hong Kong; Ada Wai-Chee Fu, Chinese University of Hong Kong; Ruifeng Liu, Chinese University of Hong Kong

Efficient Enumeration of Maximal k-Plexes

Devora Berlowitz, The Hebrew University of Jerusalem; Sara Cohen, The Hebrew University of Jerusalem; Benny Kimelfeld, Technion, Israel Institute of Technology

Divide & Conquer: I/O Efficient Depth-First Search

Zhiwei Zhang, The Chinese University of Hong Kong; Jeffrey Xu Yu, The Chinese University of Hong Kong; Lu Qin, University of Technology, Sydney; Zechao Shang, The Chinese University of Hong Kong

Index-based Optimal Algorithms for Computing Steiner Components with Maximum Connectivity

Lijun Chang, University of New South Wales; Xuemin Lin, University of New South Wales; Lu Qin, University of Technology, Sydney; Jeffrey Xu Yu, The Chinese University of Hong Kong; Wenjie Zhang, University of New South Wales

SIGMOD Research 7: Data Mining

Location: MR105

Hashtag: #sigmod15 #R7Px

Session Chair: Julia Stoyanovich (Drexel University)

COMMIT: A Scalable Approach to Mining Communication Motifs from Dynamic Networks

Saket Gururkar, IIT Madras; Sayan Ranu, IIT Madras; Balaraman Ravindran, IIT Madras

LASH: Large-Scale Sequence Mining with Hierarchies

Kaustubh Beedkar, University of Mannheim; Rainer Gemulla, University of Mannheim

Twister Tries: Approximate Hierarchical Agglomerative Clustering for Average Distance in Linear Time

Michael Cochez, University of Jyvaskyla; Hao Mou, University of Jyvaskyla

DBSCAN Revisited: Mis-Claim, Un-Fixability, and Approximation

Junhao Gan, Chinese University of Hong Kong; Yufei Tao, Chinese University of Hong Kong

The TagAdvisor: Luring the Lurkers to Review Web Items

Azade Nazi, University of Texas at Arlington; Mahashweta Das, HP Labs; Gautam Das, University of Texas at Arlington

SIGMOD Research 8: Uncertainty and Linking

Location: MR104

Hashtag: #sigmod15 #R8Px

Session Chair: Pierre Senellart (Telecom ParisTech)

Supporting Data Uncertainty in Array Databases

Liping Peng, University of Massachusetts Amherst; Yanlei Diao, University of Massachusetts Amherst

Identifying the Extent of Completeness of Query Answers over Partially Complete Databases

Simon Razniewski, Free University of Bozen-Bolzano; Flip Korn, Google Research; Werner Nutt, Free University of Bozen-Bolzano; Divesh Srivastava, AT&T Labs-Research

k-Hit Query: Top-k Query with Probabilistic Utility Function

Peng Peng, Hong Kong University of Science and Technology; Raymond Chi-Wing Wong, Hong Kong University of Science and Technology

Linking Temporal Records for Profiling Entities

Furong Li, National University of Singapore; Mong Li Lee, National University of Singapore; Wynne Hsu, National University of Singapore; Wang-Chiew Tan, UC Santa Cruz

SIGMOD Industry 2: Applications

Location: MR103

Hashtag: #sigmod15 #I2Px

Session Chair: Jignesh Patel (University of Wisconsin)

Telco Churn Prediction with Big Data

Yiqing Huang, School of Computer Science and Technology, Soochow University, Suzhou 215006, China; Fangzhou Zhu, School of Computer Science and Technology, Soochow University, Suzhou 215006, China; Mingxuan Yuan, Huawei Noah's Ark Lab, Hong Kong; Ke Deng, School of Computer Science and Information Technology RMIT University, Australia; Yanhua Li, Huawei Noah's Ark Lab, Hong Kong; Bing Ni, Huawei Noah's Ark Lab, Hong Kong; Wenyuan Dai, Huawei Noah's Ark Lab, Hong Kong; Qiang Yang, Huawei Noah's Ark Lab, Hong Kong; Jia Zeng, School of Computer Science and Technology, Soochow University, Suzhou 215006, China & Huawei Noah's Ark Lab, Hong Kong

The LDBC Social Network Benchmark: Interactive Workload

Orri Erling, OpenLink Software; Alex Averbuch, Neo Technology; Josep Larriba-Pey, Sparsity Technologies; Hassan Chafi, Oracle Labs; Andrey Gubichev, TU Munich; Arnaud Prat, Universitat Politecnica de Catalunya; Minh-Duc Pham, VU University Amsterdam; Peter Boncz, CWI

Rethinking Data-Intensive Science Using Scalable Analytics Systems

Frank Nothaft, University of California, Berkeley; Matt Massie, University of California, Berkeley; Timothy Danford, University of California, Berkeley; Zhao Zhang, University of California, Berkeley; Uri Laserson, Cloudera; Carl Yeksigian, Genomebridge; Jey Kottalam, University of California, Berkeley; Arun Ahuja, Carl Icahn School of Medicine, Mount Sinai; Jeff Hammerbacher, Carl Icahn School of Medicine, Mount Sinai;

Michael Linderman, Carl Icahn School of Medicine, Mount Sinai; Michael Franklin, University of California, Berkeley; Anthony Joseph, University of California, Berkeley; David Patterson, University of California, Berkeley

QMapper for Smart Grid: Migrating SQL-based Application to Hive

Yue Wang, Institute of Computing Technology, Chinese Academy of Sciences; Yingzhong Xu, Institute of Computing Technology, Chinese Academy of Sciences; Yue Liu, Institute of Computing Technology, Chinese Academy of Sciences; Jian Chen, Zhejiang Electric Power Corporation, China; Songlin Hu, Institute of Computing Technology, Chinese Academy of Sciences

SIGMOD Programming Contest

Location: MR106

Hashtag: #sigmod15 #pc

Session Chairs: Thomas Neumann & Alfons Kemper (Technische Universität München)

- **Introduction to the Programming Contest and presentation of this year's task - 10min**
Tobias Mühlbauer, Technische Universität München
- **Presentation of the teams' solutions - 15min per team + questions**

List of finalists (in alphabetical order):

- **Team Crisis**
Alexey Karyakin (University of Waterloo)
- **Team CStrings**
Lambros Petrou (University of Oxford)
George Koumettou (Royal Holloway University of London)
Marios Mintzis (University College London)
- **Team CUT_HERE**
Wonha Ryu (School of Computing, KAIST)
Eunjin Song (Seoul National University)
Seoung Gook Sohn (Seoul National University)
- **Team SimpleMind**
Ismail Oukid (TU Dresden)
Ingo Mueller (Karlsruhe Institute of Technology)
Iraklis Psaroudakis (École polytechnique fédérale de Lausanne)
- **Team Vika**
Thi Phuong Nhung Ngo (Free University of Bolzano)
Thai Son Nguyen (Karlsruhe Institute of Technology)

TUESDAY, 17:30 – 18:15

ACM-W Athena Lecturer Award

Location: Plenary 1

Hashtag: #sigmod15 #athenaaward

Session Chair: Magda Balazinska (University of Washington)

Three Favorite Results

Jennifer Widom, Stanford University

TUESDAY, 18:25 – 18:55

SIGMOD Business Meeting

Location: Plenary 1

Session Chair: Kossmann Donald (ETH Zürich)

TUESDAY, 18:30 – 19:30

Careers in Industry

Location: MR106

Hashtag: #sigmod15 #career

Session Chair: Feifei Li (University of Utah)

WEDNESDAY, 8:20 – 10:00

SIGMOD Keynote Talk 2

Location: Plenary 1

Hashtag: #sigmod15 #keynote2

Session Chair: Zack Ives (University of Pennsylvania)

The Power Behind the Throne: Information Integration in the Age of Data-Driven Discovery

Laura M. Haas, IBM Research Accelerated Discovery Lab

WEDNESDAY, 10:30 – 12:10

PODS Session 8: Streaming (ending at 12:00)

Location: Plenary 1

Hashtag: #pods15 #R8Px

Session Chair: Graham Cormode (University of Warwick)

External Memory Stream Sampling

Xiaocheng Hu, the Chinese University of Hong Kong; Miao Qiao, the Chinese University of Hong Kong; Yufei Tao, the Chinese University of Hong Kong

Vertex and Hyperedge Connectivity in Dynamic Graph Streams

Sudipto Guha, University of Pennsylvania; Andrew McGregor, University of Massachusetts; David Tench, University of Massachusetts

Fast and Near-Optimal Algorithms for Approximating Distributions by Histograms

Jayadev Acharya, MIT; Ilias Diakonikolas, University of Edinburgh; Chinmay Hegde, MIT; Jerry Li, MIT; Ludwig Schmidt, MIT

SIGMOD Research 9: Transactional Architectures

Location: MR103

Hashtag: #sigmod15 #R9Px

Session Chair: Feifei Li (University of Utah)

On the Design and Scalability of Distributed Shared-Data Databases

Simon Loesing, ETH Zurich; Markus Pilman, ETH Zurich; Thomas Etter, ETH Zurich; Donald Kossmann, Microsoft Research

Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems

Thomas Neumann, Technische Universität München; Tobias Mühlbauer, Technische Universität München; Alfons Kemper, Technische Universität München

FOEDUS: OLTP Engine for a Thousand Cores and NVRAM

Hideaki Kimura, HP Labs

Let's Talk About Storage & Recovery Methods for Non-Volatile Memory Database Systems

Joy Arulraj, Carnegie Mellon University; Andrew Pavlo, Carnegie Mellon University; Subramanya Dullloor, Intel Labs

SIGMOD Research 10: Privacy

Location: MR104

Hashtag: #sigmod15 #R10Px

Session Chair: Wang-Chiew Tan (University of California, Santa Cruz)

Private Release of Graph Statistics using Ladder Functions

Jun Zhang, Nanyang Technological University; Graham Cormode, University of Warwick; Cecilia Procopiuc, Google Inc.; Divesh Srivastava, AT&T Labs - Research; Xiaokui Xiao, Nanyang Technological University

Bayesian Differential Privacy on Correlated Data

Bin Yang, Rakuten Inc.; Issei Sato, The University of Tokyo; Hiroshi Nakagawa, The University of Tokyo

Modular Order-Preserving Encryption, Revisited

Charalampos Mavroforakis, Boston University; Nathan Chenette, Rose-Hulman Institute of Technology; Adam O'Neill, Georgetown University and NIST; George Kollios, Boston University; Ran Canetti, Boston University

Chiaroscuro: Transparency and Privacy for Massive Personal Time-Series Clustering

Tristan Allard, IRISA & Univ. Rennes 1; Georges Hébrail, EDF R&D; Florent Masseglia, Inria & Lirmm, Univ. Montpellier; Esther Pacitti, Inria & Lirmm, Univ. Montpellier

SIGMOD Research 11: Streams

Location: MR101&102

Hashtag: #sigmod15 #R11Px

Session Chair: Jorge Quiané-Ruiz (Qatar Computing Research Institute)

Persistent Data Sketching

Zhewei Wei, Renmin University of China; Ge Luo, The Hong Kong University of Science and Technology; Ke Yi, The Hong Kong University of Science and Technology; Xiaoyong Du, Renmin University of China; Ji-Rong Wen, Renmin University of China

Scalable Distributed Stream Join Processing

Qian Lin, National University of Singapore; Beng Chin Ooi, National University of Singapore; Zhengkui Wang, National University of Singapore; Cui Yu, Monmouth University

SCREEN: Stream Data Cleaning under Speed Constraints

Shaouxu Song, Tsinghua University; Aoqian Zhang, Tsinghua University; Jianmin Wang, Tsinghua University; Philip Yu, University of Illinois at Chicago

Location-Aware Pub/Sub System: When Continuous Moving Queries Meet Dynamic Event Streams

Long Guo, National University of Singapore; Dongxiang Zhang, National University of Singapore; Guoliang Li, Tsinghua University; Kian-Lee Tan, National University of Singapore; Zhifeng Bao, RMIT University

SIGMOD Demonstrations, Group A

Location: MR106

Hashtag: #sigmod15 #demoAx

CE-Storm: Confidential Elastic Processing of Data Streams

Nick Katsipoulakis, University of Pittsburgh; Cory Thoma, University of Pittsburgh; Eric Gratta, University of Pittsburgh; Alexandros Labrinidis, University of Pittsburgh; Adam Lee, University of Pittsburgh; Panos Chrysanthis, University of Pittsburgh

IReS: Intelligent, Multi-Engine Resource Scheduler for Big Data Analytics Workflows

Katerina Doka, National Technical University of Athens; Nikolaos Papailiou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Christos Mantas, National Technical University of Athens; Nectarios Koziris, National Technical University of Athens

Exploratory Keyword Search with Interactive Input

Zhifeng Bao, RMIT University; Yong Zeng, National University of Singapore; H. V. Jagadish, School of Computing; Tok Wang Ling, National University of Singapore

QE3D: Interactive Visualization and Exploration of Complex, Distributed Query Plans

Daniel Scheibli, SAP SE; Christian Dinse, SAP SE; Alexander Boehm, SAP SE

DataXFormer: An Interactive Data Transformation Tool

John Morcos, University of Waterloo; Ziawasch Abedjan, CSAIL MIT; Ihab Francis Ilyas, University of Waterloo; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Michael Stonebraker, MIT CSAIL

Quality-Driven Continuous Query Execution over Out-of-Order Data Streams

Yuanzhen Ji, SAP SE; Hongjin Zhou, SAP SE; Zbigniew Jerzak, SAP SE; Anisoara Nica, SAP SE; Gregor Hackenbroich, SAP SE; Christof Fetzer, Technische Universitaet Dresden

MoDisSENSE: A Distributed Spatio-Temporal and Textual Processing Platform for Social Networking Services

Ioannis Mytilinis, National Technical University of Athens; Ioannis Giannakopoulos, National Technical University of Athens; Ioannis Konstantinou, National Technical University of Athens; Katerina Doka, National Technical University of Athens; Dimitrios Tsitsigkos, IMIS, RC Athena; Manolis Terrovitis, IMIS, RC Athena; Lampros Giampouras, Athens Technology Center S.A; Nectarios Koziris, National Technical University of Athens

DocRicher: An Automatic Annotation System for Text Documents Using Social Media

Qiang Hu, SeSaMe Centre; Qi Liu, School of Computing; Xiaoli Wang, SeSaMe Centre; Anthony Tung, School of Computing; Shubham Goyal, SeSaMe Centre; Jisong Yang, SeSaMe Centre

A Demonstration of Rubato DB: A Highly Scalable NewSQL Database System for OLTP and Big Data Applications

Li-Yan Yuan, University of Alberta; Lengdong Wu, University of Alberta; Jia-Huai You, Computing Science; Yan Chi, Shanghai Shifang Software Ltd

G-OLA: Generalized On-Line Aggregation for Interactive Analysis on Big Data

Kai Zeng, University of California, Berkeley; Sameer Agarwal, Databricks Inc.; Ankur Dave, University of California, Berkeley; Michael Armbrust, Databricks Inc.; Ion Stoica, University of California, Berkeley

SIGMOD Tutorial 2.1: Mining and Forecasting of Big Time-series Data

Location: MR105

Hashtag: #sigmod15 #tutorial2

Presenters: Yasushi Sakurai, Kumamoto University; Yasuko Matsubara, Kumamoto University; Christos Faloutsos, Carnegie Mellon University

WEDNESDAY, 12:10–13:20

Lunch and Posters

Location: Level 1 Foyer

WEDNESDAY, 13:20 – 14:50

SIGMOD New Researcher Symposium

Location: Plenary 1

Hashtag: #sigmod15 #nrs

Session Chairs: Alexandra Meliou (University of Massachusetts Amherst); Julia Stoyanovich (Drexel University)

PODS Session 9: Algorithms and Data Structures (ending at 16:50)

Location: Plenary 1

Hashtag: #pods15 #R9Px

Session Chair: Benny Kimelfeld (LogicBlox)

On Top-k Range Reporting in 2D Space

Rahul Saladi, University of Minnesota; Yufei Tao, Chinese University of Hong Kong.

Dynamic Data Structures for Document Collections and Graphs

Ian Munro, University of Waterloo; Yakov Nekrich, University of Waterloo; Jeff Vitter, University of Kansas.

Join Dependency Testing, Loomis-Whitney Join, and Triangle Enumeration

Xiaocheng Hu, Chinese University of Hong Kong; Miao Qiao, Chinese University of Hong Kong; Yufei Tao, Chinese University of Hong Kong

SIGMOD Research 12: Spatial data

Location: MR103

Hashtag: #sigmod15 #R12Px

Session Chair: Cyrus Shahabi (University of Southern California)

Optimal Spatial Dominance: An effective search of Nearest Neighbor Candidates

Xiaoyang Wang, The University Of New South Wales; Ying Zhang, University of Technology, Sydney; Wenjie Zhang, The University Of New South Wales; Xuemin Lin, The University Of New South Wales; Muhammad Aamir Cheema, Monash University

THERMAL-JOIN: A Scalable Spatial Join for Dynamic Workloads

Farhan Tauheed, Oracle Labs; Thomas Heinis, Imperial College; Anastasia Ailamaki, Ecole Polytechnique Federale de Lausanne

Indexing Metric Uncertain Data for Range Queries

Lu Chen, Zhejiang University; Yunjun Gao, Zhejiang University; Xinhan Li, Zhejiang University; Christian S. Jensen, Aalborg University; Gang Chen, Zhejiang University; Baihua Zheng, Singapore Management University

Efficient Route Planning on Public Transportation Networks: A Labelling Approach

Sibo Wang, Nanyang Technological University; Wenqing Lin, A*STAR; Yi Yang, Fudan University; Xiaokui Xiao, Nanyang Technological University; Shuigeng Zhou, Fudan University

SIGMOD Research 13: Crowdsourcing

Location: MR104

Hashtag: #sigmod15 #R13Px

Session Chair: Dimitrios Tsoumakos (National Technical University of Athens)

The Importance of Being Expert: Efficient Max-Finding in Crowdsourcing

Aris Anagnostopoulos, Sapienza University of Rome; Luca Becchetti, Sapienza University of Rome; Adriano Fazzone, Sapienza University of Rome; Ida Mele, Max Planck Institute for Informatics; Matteo Riondato, Stanford University

Minimizing Efforts in Validating Crowd Answers

Quoc Viet Hung Nguyen, EPFL; Chi Thang Duong, EPFL; Matthias Weidlich, Imperial College London; Karl Aberer, EPFL

iCrowd: An Adaptive Crowdsourcing Framework

Ju Fan, School of Computing, National University of Singapore; Guoliang Li, Department of Computer Science, Tsinghua University; Beng Chin Ooi, School of Computing, National University of Singapore; Kian-lee Tan, School of Computing, National University of Singapore; Jianhua Feng, Department of Computer Science, Tsinghua University

QASCA: A Quality-Aware Task Assignment System for Crowdsourcing Applications

Yudian Zheng, The University of Hong Kong; Jiannan Wang, UC Berkeley; Guoliang Li, Tsinghua University; Reynold Cheng, The University of Hong Kong; Jianhua Feng, Tsinghua University

tDP: An Optimal-Latency Budget Allocation strategy for Crowdsourced MAXIMUM operations

Vasilis Verraios, Stanford University; Peter Lofgren, Stanford University; Hector Garcia-Molina, Stanford University

SIGMOD Demonstrations, Group B

Location: MR106

Hashtag: #sigmod15 #demoBx

Thrifty: Offering Parallel Database as a Service using the Shared-Process Approach

Petrie Wong, The University of Hong Kong; Zhian He, The Hong Kong Polytechnic University; Ziqiang Feng, The Hong Kong Polytechnic University; Wenjian Xu, The Hong Kong Polytechnic University; Eric Lo, The Hong Kong Polytechnic University

BenchPress: Dynamic Workload Control in the OLTP-Bench Testbed

Dana Van Aken, Carnegie Mellon University; Djellel Difallah, University of Fribourg; Andrew Pavlo, Carnegie Mellon University; Carlo Curino, Microsoft Corporation; Philippe Cudre-Mauroux, University of Fribourg

Demonstrating Data Near Here - Scientific Data Search

V.M. Megler, Portland State University; David Maier, Portland State University

Slider: an Efficient Incremental Reasoner

Jules Chevalier, Université Jean Monnet; Julien Subercaze, Université Jean Monnet; Christophe Gravier, Université Jean Monnet; Frédérique Laforest, Université Jean Monnet

WANalytics: Geo-Distributed Analytics for a Data Intensive World

Ashish Vulimiri, UIUC; Carlo Curino, Microsoft; Philip Godfrey, UIUC; Thomas Jungblut, Microsoft; Konstantinos Karanasos, Microsoft; Jitendra Padhye, Microsoft; George Varghese, Microsoft

FTT: a System for Finding and Tracking Tourists in Public Transport Services

Huayu Wu, Institute for Infocomm Research; Jo-Anne Tan, Institute for Infocomm Research; Wee Siong Ng, Institute for Infocomm Research; Mingqiang Xue, Institute for Infocomm Research; Wei Chen, Institute for Infocomm Research

SharkDB: An In-Memory Storage System for Massive Trajectory Data

Haozhou Wang, The University of Queensland; Kai Zheng, The University of Queensland; Xiaofang Zhou, The University of Queensland; Shazia Sadiq, The University of Queensland

Ringo: Interactive Graph Analytics on Big-Memory Machines

Yonathan Perez, Stanford University; Rok Susic, Stanford University; Arijit Banerjee, Stanford University; Rohan Puttagunta, Stanford University; Martin Raison, Stanford University; Pararth Shah, Stanford University; Jure Leskovec, Stanford University

STORM: Spatio-Temporal Online Reasoning and Management of Large Spatio-Temporal Data

Robert Christensen, University of Utah; Feifei Li, University of Utah; Lu Wang, Hong Kong University of Science and Technology; Ke Yi, Hong Kong; Jun Tang, University of Utah; Natalee Villa, University of Utah

PAXQuery: Parallel Analytical XML Processing

Jesús Camacho-Rodríguez, Hortonworks Inc.; Dario Colazzo, Université Paris-Dauphine; Ioana Manolescu, INRIA & Université Paris-Sud; Juan Naranjo, INRIA & Université Paris-Sud

SIGMOD Tutorial 2.2: Mining and Forecasting of Big Time-series Data (cont.)

Location: MR105

Hashtag: #sigmod15 #tutorial2

Presenters: Yasushi Sakurai, Kumamoto University; Yasuko Matsubara, Kumamoto University; Christos Faloutsos, Carnegie Mellon University

PODS Session 10: Weights and Similarity (ending at 18:40)

Location: Plenary 1

Hashtag: #pods15 #R9Px

Session Chair: Werner Nutt (Free University of Bozen-Bolzano)

Private Multiplicative Weights Beyond Linear Queries

Jonathan Ullman, Columbia University

Symmetric Weighted First-Order Model Counting

Paul Beame, University of Washington; Guy Van den Broeck, KU Leuven; Eric Gribkoff, University of Washington; Dan Suciu, University of Washington

Smooth Tradeoffs between Insert and Query Complexity in Nearest Neighbor Search

Michael Kapralov, IBM Watson

SIGMOD Research 14: Indexing & Performance

Location: MR101&102

Hashtag: #sigmod15 #R14Px

Session Chair: Yufei Tao (The Chinese University of Hong Kong)

Cache-Efficient Aggregation: Hashing is Sorting

Ingo Müller, Karlsruhe Institute of Technology / SAP SE; Peter Sanders, Karlsruhe Institute of Technology; Arnaud Lacurie, SAP SE; Wolfgang Lehner, Dresden University of Technology; Franz Färber, SAP SE

Efficient Similarity Join and Search on Multi-Attribute Data

Guoliang Li, Tsinghua University; Jian He, Tsinghua University; Dong Deng, Tsinghua University; Jian Li, Tsinghua University

Holistic Indexing in Main-memory Column-stores

Eleni Petraki, CWI; Stratos Idreos, Harvard University; Stefan Manegold, CWI

CliffGuard: A Principled Framework for Finding Robust Database Designs

Barzan Mozafari, University of Michigan, Ann Arbor; Eugene Goh, University of Michigan, Ann Arbor; Dong Yoon, University of Michigan, Ann Arbor

Exploiting Correlations for Expensive Predicate Evaluation

Manas Joglekar, Stanford University; Hector Garcia-Molina, Stanford University; Aditya Parameswaran, University of Illinois at Urbana-Champaign; Christopher Ré, Stanford University

SIGMOD Research 15: Data Cleaning

Location: MR103

Hashtag: #sigmod15 #R15Px

Session Chair: Mohamad Sharaf (University of Queensland)

Query-Oriented Data Cleaning with Oracles

Moria Bergman, Tel Aviv University; Tova Milo, Tel Aviv University; Slava Novgorodov, Tel Aviv University; Wang-Chiew Tan, UC Santa Cruz

BigDancing: A System for Big Data Cleansing

Zuhair Khayyat, King Abdullah University of Science and Technology (KAUST); Ihab Ilyas, University of Waterloo; Alekh Jindal, MIT; Samuel Madden, MIT; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Jorge-Arnelo Quiané-Ruiz, Qatar Computing Research Institute; Nan Tang, Qatar Computing Research Institute; Si Yin, Qatar Computing Research Institute

Data X-Ray: A diagnostic tool for data errors

Xiaolan Wang, University of Massachusetts, Amherst; Xin Dong, Google Inc.; Alexandra Meliou, University of Massachusetts, Amherst

KATARA: A Data Cleaning System Powered by Knowledge Bases and Crowdsourcing

Xu Chu, University of Waterloo; John Morcos, University of Waterloo; Ihab Ilyas, University of Waterloo; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Nan Tang, Qatar Computing Research Institute; Yin Ye, Qatar Computing Research Institute

Crowd-based Deduplication: an Adaptive Approach

Sibo Wang, Nanyang Technological University; Xiaokui Xiao, Nanyang Technological University; Chun-Hee Lee, Nanyang Technological University

SIGMOD Research 16: Transactions

Location: MR104

Hashtag: #sigmod15 #R16Px

Session Chair: Ashraf Aboulnaga (University of Waterloo)

Minimizing Commit Latency of Transactions in Geo-Replicated Data Stores

Faisal Nawab, UC Santa Barbara; Vaibhav Arora, UC Santa Barbara; Divyakant Agrawal, UC Santa Barbara; Amr El Abbadi, UC Santa Barbara

Optimizing Optimistic Concurrency Control for Tree-Structured, Log-Structured Databases

Philip Bernstein, Microsoft Research; Sudipto Das, Microsoft Research; Bailu Ding, Cornell University; Markus Pilman, ETH Zurich

The Homeostasis Protocol: Avoiding Transaction Coordination Through Program Analysis

Sudip Roy, Google Research & Cornell University; Lucja Kot, Cornell University; Gabriel Bender, Cornell University; Bailu Ding, Cornell University; Hossein Hojjat, Cornell University; Christoph Koch, EPFL; Nate Foster, Cornell University; Johannes Gehrke, Microsoft Corp. & Cornell University

Feral Concurrency Control: An Empirical Investigation of Modern Application Integrity

Peter Bailis, UC Berkeley; Alan Fekete, University of Sydney; Michael Franklin, UC Berkeley; Ali Ghodsi, UC Berkeley; Joseph Hellerstein, UC Berkeley; Ion Stoica, UC Berkeley

SIGMOD Industry 3: Novel Systems

Location: MR105

Hashtag: #sigmod15 #I3Px

Session Chair: Theo Vassilakis (Metanautix)

REEF: Retainable Evaluator Execution Framework

Markus Weimer, Microsoft; Yingda Chen, Microsoft; Byung-Gon Chun, Seoul National University; Tyson Condie, UCLA; Carlo Curino, Microsoft; Chris Douglas, Microsoft; Yunseong Lee, Seoul National University; Tony Majestro, Microsoft; Dahlia Malkhi, VMware; Sergiy Matusevych, Microsoft; Brandon Myers, University of Washington; Shravan Narayanamurthy, Microsoft; Raghu Ramakrishnan, Microsoft; Sriram Rao, Microsoft; Russel Sears, PureStorage; Beysim Sezgin, Microsoft; Julia Wang, Microsoft

Apache Tez: A Unifying Framework for Modeling and Building Data Processing Applications

Bikas Saha, Hortonworks; Hitesh Shah, Hortonworks; Siddharth Seth, Hortonworks; Gopal Vijayaraghavan, Hortonworks; Arun Murthy, Hortonworks; Carlo Curino, Microsoft

Design and Implementation of the LogicBlox System

Molham Aref, LogicBlox; Balder ten Cate, LogicBlox; Todd Green, LogicBlox; Benny Kimelfeld, LogicBlox; Dan Olteanu, LogicBlox; Emir Pasalic, LogicBlox; Todd Veldhuizen, LogicBlox; Geoffrey Washburn, LogicBlox

Spark SQL: Relational Data Processing in Spark

Michael Armbrust, Databricks; Reynold Xin, Databricks; Cheng Lian, Databricks; Yin Huai, Databricks; Davies Liu, Databricks; Joseph Bradley, Databricks; Xiangrui Meng, Databricks; Tomer Kaftan, UC Berkeley; Michael Franklin, UC Berkeley; Ali Ghodsi, Databricks; Matei Zaharia, Databricks

SIGMOD Demonstrations, Group C

Location: MR106

Hashtag: #sigmod15 #demoCx

Graft: A Debugging Tool For Apache Giraph

Semih Salihoglu, Stanford University; Jaeho Shin, Stanford University; Vikesh Khanna, Stanford University; Ba Quan Truong, Stanford University; Jennifer Widom, Stanford University

Even Metadata is Getting Big: Annotation Summarization using InsightNotes

Dongqing Xiao, Worcester Polytechnic Institute; Amir Bashlari, Worcester Polytechnic Institute; Tyler Menard, Worcester Polytechnic Institute; Mohamed Eltabakh, Worcester Polytechnic Institute

StoryPivot: Comparing and Contrasting Story Evolution

Anja Gruenheid, ETH Zurich; Donald Kossmann, ETH Zurich; Theodoros Rekatsinas, University of Maryland; Divesh Srivastava, AT&T Labs-Research

The Flatter, the Better – Query Compilation Based on the Flattening Transformation

Alexander Ulrich, Universität Tübingen; Torsten Grust, Universität Tübingen

D2WORM: A Management Infrastructure for Distributed Data-centric Workflows

Martin Jergler, Technische Universität München; Mohammad Sadoghi, IBM; Hans-Arno Jacobsen, Technische Universität München

NL2CM: A Natural Language Interface to Crowd Mining

Yael Amsterdamer, Tel Aviv University; Anna Kukliansky, Tel Aviv University; Tova Milo, Tel Aviv University

Optimistic Recovery for Iterative Dataflows in Action

Sergey Dudoladov, Technische Universität Berlin; Chen Xu, Technische Universität Berlin; Sebastian Schelter, Technische Universität Berlin; Asterios Katsifodimos, Technische Universität Berlin; Stephan Ewen, Data Artisans GmbH; Kostas Tzoumas, Data Artisans GmbH; Volker Markl, Technische Universität Berlin

A Secure Search Engine for the Personal Cloud

Saliha Lallali, INRIA Rocquencourt; Nicolas Ancaux, INRIA Rocquencourt; Iulian Sandu Popa, INRIA Rocquencourt; Philippe Pucheral, INRIA Rocquencourt

A SQL Debugger Built from Spare Parts - Turning a SQL:1999 Database System into Its Own Debugger

Benjamin Dietrich, Universität Tübingen; Torsten Grust, Universität Tübingen

Just can't get enough - Synthesizing Big Data

Tilmann Rabl, University of Toronto; Manuel Danisch, bankmark; Michael Frank, bankmark; Sebastian Schindler, bankmark; Hans-Arno Jacobsen, University of Toronto

WEDNESDAY, 19:15 – 22:00

Conference Banquet

Location: Melbourne Town Hall

THURSDAY, 8:30 – 10:00

SIGMOD Awards Talks

Location: Plenary 1

Hashtag: #sigmod15 #award

Session Chair: Elisa Bertino (Purdue University)

THURSDAY, 10:30 – 12:10

SIGMOD Research 17: Hardware-Aware Query Processing

Location: MR103

Hashtag: #sigmod15 #R17Px

Session Chair: Spyros Blanas (Ohio State University)

Rack-Scale In-Memory Join Processing using RDMA

Claude Barthels, ETH Zurich; Simon Loesing, ETH Zurich; Gustavo Alonso, ETH Zurich; Donald Kossmann, ETH Zurich

Self-Tuning, GPU-Accelerated Kernel Density Models for Multidimensional Selectivity Estimation

Max Heimerl, Technische Universität Berlin; Martin Kiefer, Technische Universität Berlin; Volker Markl, Technische Universität Berlin

Rethinking SIMD Vectorization for In-Memory Databases

Orestis Polychroniou, Columbia University; Arun Raghavan, Oracle Corporation; Kenneth Ross, Columbia University

A Padded Encoding Scheme to Accelerate Scans by Leveraging Skew

Yinan Li, University of Wisconsin-Madison; Craig Chasseur, University of Wisconsin-Madison; Jignesh Patel, University of Wisconsin-Madison

SIGMOD Research 18: Graph Propagation, Influence, Mining

Location: MR101&102

Hashtag: #sigmod15 #R18Px

Session Chair: Panagiotis Karras (Skoltech)

GetReal: Towards Realistic Selection of Influence Maximization Strategies in Competitive Networks

Hui Li, Xidian University; Sourav S Bhowmick, Nanyang Technological University; Jiangtao Cui, Xidian University; Yunjun Gao, Zhejiang University

Influence Maximization in Near-Linear Time: A Martingale Approach

Youze Tang, Nanyang Technological University; Yanchen Shi, Nanyang Technological University; Xiaokui Xiao, Nanyang Technological University

Community Level Diffusion Extraction

Zhiting Hu, Peking University & Carnegie Mellon University; Junjie Yao, East China Normal University; Bin Cui, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University; Eric Xing, Language Technologies Institute, Carnegie Mellon University

BEAR: Block Elimination Approach for Random Walk with Restart on Large Graphs

Kijung Shin, Seoul National University; Jinhong Jung, KAIST; Sael Lee, The State University of New York (SUNY) Korea; U Kang, KAIST

The Minimum Wiener Connector Problem

Natali Ruchansky, Boston University; Francesco Bonchi, Yahoo Labs; David García-Soriano, Yahoo Labs; Francesco Gullo, Yahoo Labs; Nicolas Kourtellis, Yahoo Labs

SIGMOD Research 19: Social Networks

Location: MR104

Hashtag: #sigmod15 #R19Px

Session Chair: Chengfei Liu (Swinburne University of Technology)

From Group Recommendations to Group Formation

Senjuti Basu Roy, University of Washington Tacoma; Laks Lakshmanan, University of British Columbia; Rui Liu, University of Washington Tacoma

Real-Time Multi-Criteria Social Graph Partitioning: A Game Theoretic Approach

Nikos Armenatzoglou, Hong Kong University of Science and Technology; Huy Pham, University of Southern California; Vasilis Ntranos, University of Southern California; Dimitris Papadias, Hong Kong University of Science and Technology; Cyrus Shahabi, University of Southern California

Utility-Aware Social Event-Participant Planning

Jieying She, The Hong Kong University of Science and Technology; Yongxin Tong, Beihang University; Lei Chen, The Hong Kong University of Science and Technology

Online Video Recommendation in Sharing Community

Xiangmin Zhou, RMIT University; Lei Chen, Hong Kong University of Science and Technology; Yanchun Zhang, Victoria University; Longbing Cao, University of Technology, Sydney; Guangyan Huang, Deakin University; Chen Wang, CSIRO

SIGMOD Industry 4: Performance

Location: MR105

Hashtag: #sigmod15 #4Px

Session Chair: Yanlei Diao (University of Massachusetts, Amherst)

Large-scale Predictive Analytics in Vertica: Fast Data Transfer, Distributed Model Creation, and In-database Prediction

Shreya Prasad, HP Vertica; Arash Fard, HP Vertica; Vishrut Gupta, HP Vertica; Jorge Martinez, HP Vertica; Jeff LeFevre, HP Vertica; Vincent Xu, HP Vertica; Meichun Hsu, HP Vertica; Indrajit Roy, HP Labs

Oracle Workload Intelligence

Quoc Trung Tran, Oracle USA; Konstantinos Morfonios, Oracle USA; Neoklis Polyzotis, Google, Inc.

Purity: Building Fast, Highly-Available Enterprise Flash Storage from Commodity Components

John Colgrove, Pure Storage; John Davis, Pure Storage; John Hayes, Pure Storage; Ethan Miller, Pure Storage; Cary Sandvig, Pure Storage; Russell Sears, Pure Storage; Ari Tamches, Pure Storage; Neil Vachharajani, Pure Storage; Feng Wang, Pure Storage

On Improving User Response Times in Tableau

Pawel Terlecki, Tableau Software; Fei Xu, Tableau Software; Marianne Shaw, Tableau Software; Valeri Kim, Tableau Software; Richard Wesley, Tableau Software

SIGMOD Demonstrations, Group B

Location: MR106

Hashtag: #sigmod15 #demoBx

Thrifty: Offering Parallel Database as a Service using the Shared-Process Approach

Petrie Wong, The University of Hong Kong; Zhian He, The Hong Kong Polytechnic University; Ziqiang Feng, The Hong Kong Polytechnic University; Wenjian Xu, The Hong Kong Polytechnic University; Eric Lo, The Hong Kong Polytechnic University

BenchPress: Dynamic Workload Control in the OLTP-Bench Testbed

Dana Van Aken, Carnegie Mellon University; Djellel Difallah, University of Fribourg; Andrew Pavlo, Carnegie Mellon University; Carlo Curino, Microsoft Corporation; Philippe Cudre-Mauroux, University of Fribourg

Demonstrating Data Near Here - Scientific Data Search

V.M. Megler, Portland State University; David Maier, Portland State University

Slider: an Efficient Incremental Reasoner

Jules Chevalier, Université Jean Monnet; Julien Subercaze, Université Jean Monnet; Christophe Gravier, Université Jean Monnet; Frédérique Laforest, Université Jean Monnet

WANalytics: Geo-Distributed Analytics for a Data Intensive World

Ashish Vulimiri, UIUC; Carlo Curino, Microsoft; Philip Godfrey, UIUC; Thomas Jungblut, Microsoft; Konstantinos Karanasos, Microsoft; Jitendra Padhye, Microsoft; George Varghese, Microsoft

FTT: a System for Finding and Tracking Tourists in Public Transport Services

Huayu Wu, Institute for Infocomm Research; Jo-Anne Tan, Institute for Infocomm Research; Wee Siong Ng, Institute for Infocomm Research; Mingqiang Xue, Institute for Infocomm Research; Wei Chen, Institute for Infocomm Research

SharkDB: An In-Memory Storage System for Massive Trajectory Data

Haozhou Wang, The University of Queensland; Kai Zheng, The University of Queensland; Xiaofang Zhou, The University of Queensland; Shazia Sadiq, The University of Queensland

Ringo: Interactive Graph Analytics on Big-Memory Machines

Yonathan Perez, Stanford University; Rok Susic, Stanford University; Arijit Banerjee, Stanford University; Rohan Puttagunta, Stanford University; Martin Raison, Stanford University; Pararth Shah, Stanford University; Jure Leskovec, Stanford University

STORM: Spatio-Temporal Online Reasoning and Management of Large Spatio-Temporal Data

Robert Christensen, University of Utah; Feifei Li, University of Utah; Lu Wang, Hong Kong University of Science and Technology; Ke Yi, Hong Kong; Jun Tang, University of Utah; Natalee Villa, University of Utah

PAXQuery: Parallel Analytical XML Processing

Jesús Camacho-Rodríguez, Hortonworks Inc.; Dario Colazzo, Université Paris-Dauphine; Ioana Manolescu, INRIA & Université Paris-Sud; Juan Naranjo, INRIA & Université Paris-Sud

SIGMOD Tutorial 3.1: Data Management in Non-Volatile Memory

Location: Plenary 1

Hashtag: #sigmod15 #tutorial3

Presenters: Stratis D. Viglas, University of Edinburgh

THURSDAY, 12:10–13:30

Lunch and Posters

Location: Level 1 Foyer

THURSDAY, 13:30 – 15:10

SIGMOD Research 20: Information Extraction and Record Linking

Location: MR103

Hashtag: #sigmod15 #R20Px

Session Chair: Guoliang Li (Tsinghua University)

TEGRA: Table Extraction by Global Record Alignment

Xu Chu, University of Waterloo; Yeye He, Microsoft Research; Kaushik Chakrabarti, Microsoft Research; Kris Ganjam, Microsoft Research

Mining Quality Phrases from Massive Text Corpora

Jialu Liu, University of Illinois at Urbana-Champaign; Jingbo Shang, University of Illinois at Urbana-Champaign; Chi Wang, Microsoft Research; Xiang Ren, University of Illinois at Urbana-Champaign; Jiawei Han, University of Illinois at Urbana-Champaign

Mining Subjective Properties on the Web

Immanuel Trummer, EPFL; Alon Halevy, Google, Inc.; Hongrae Lee, Google, Inc.; Sunita Sarawagi, Google, Inc.; Rahul Gupta, Google, Inc.

Microblog Entity Linking with Social Temporal Context

Wen Hua, The University of Queensland; Kai Zheng, The University of Queensland; Xiaofang Zhou, The University of Queensland

SIGMOD Research 21: RDF and SPARQL

Location: MR101&102

Hashtag: #sigmod15 #R21Px

Session Chair: Nikos Mamoulis (The University of Hong Kong)

Graph-Aware, Workload-Adaptive SPARQL Query Caching

Nikolaos Papailiou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Panagiotis Karras, Skolkovo Institute of Science and Technology (Skoltech); Nectarios Koziris, National Technical University of Athens

Left Bit Right: For SPARQL Join Queries with OPTIONAL Patterns (Left-outer-joins)

Medha Atre, Unaffiliated

How to Build Templates for RDF Question/Answering -- An Uncertain Graph Similarity Join Approach

Weiguo Zheng, Peking University; Lei Zou, Peking University; Xiang Lian, University of Texas-Pan American; Jeffrey Xu Yu, The Chinese University of Hong Kong; Shaoxu Song, Tsinghua University; Dongyan Zhao, Peking University

RBench: Application-Specific RDF Benchmarking

Shi Qiao, Case Western Reserve University; Z. Meral Özsoyoğlu, Case Western Reserve University

ALEX: Automatic Link Exploration in Linked Data

Ahmed El-Roby, University of Waterloo; Ashraf Aboulnaga, Qatar Computing Research Institute

SIGMOD Research 22: Time Series & Graph Processing

Location: MR104

Hashtag: #sigmod15 #R22Px

Session Chair: Alan Fekete (University of Sydney)

k-Shape: Efficient and Accurate Clustering of Time Series

John Paparrizos, Columbia University; Luis Gravano, Columbia University

SMILer: A Semi-Lazy Time Series Prediction System for Sensors

Jingbo Zhou, National University of Singapore; Anthony Tung, National University of Singapore

SQLGraph: An Efficient Relational-Based Property Graph Store

Wen Sun, IBM Research - China; Achille Fokoue, IBM Watson Research Center; Kavitha Srinivas, IBM Watson Research Center; Anastasios Kementsietsidis, Google Inc.; Gang Hu, IBM Research - China; Guotong Xie, IBM Research - China

Updating Graph Indices with a One-Pass Algorithm

Dayu Yuan, Google; Prasenjit Mitra, QCRI; Huiwen Yu, Google; C. Lee Giles, Penn State University

SIGMOD Industry 5: Usability

Location: MR105

Hashtag: #sigmod15 #I5Px

Session Chair: Magda Balazinska (University of Washington)

Amazon Redshift and the Case for Simpler Data Warehouses

Anurag Gupta, Amazon Web Services; Deepak Agarwal, Amazon Web Services; Derek Tan, Amazon Web Services; Jakub Kulesza, Amazon Web Services; Rahul Pathak, Amazon Web Services; Stefano Stefani, Amazon Web Services; Vidhya Srinivasan, Amazon

ShareInsights - An Unified Approach to Full-stack Data Processing

MUKUND DESHPANDE, Persistent Systems Ltd.; DHRUVA RAY, Persistent Systems Ltd.; SAMEER DIXIT, Persistent Systems Ltd.; AVADHOOT AGASTI, Persistent Systems Ltd.

SIGMOD Demonstrations, Group C

Location: MR106

Hashtag: #sigmod15 #demoCx

Graft: A Debugging Tool For Apache Giraph

Semih Salihoglu, Stanford University; Jaeho Shin, Stanford University; Vikesh Khanna, Stanford University; Ba Quan Truong, Stanford University; Jennifer Widom, Stanford University

Even Metadata is Getting Big: Annotation Summarization using InsightNotes

Dongqing Xiao, Worcester Polytechnic Institute; Armir Bashllari, Worcester Polytechnic Institute; Tyler Menard, Worcester Polytechnic Institute; Mohamed Eltabakh, Worcester Polytechnic Institute

StoryPivot: Comparing and Contrasting Story Evolution

Anja Gruenheid, ETH Zurich; Donald Kossmann, ETH Zurich; Theodoros Rekatsinas, University of Maryland; Divesh Srivastava, AT&T Labs-Research

The Flatter, the Better -- Query Compilation Based on the Flattening Transformation

Alexander Ulrich, Universität Tübingen; Torsten Grust, Universität Tübingen

D2WORM: A Management Infrastructure for Distributed Data-centric Workflows

Martin Jergler, Technische Universität München; Mohammad Sadoghi, IBM; Hans-Arno Jacobsen, Technische Universität München

NL2CM: A Natural Language Interface to Crowd Mining

Yael Amsterdamer, Tel Aviv University; Anna Kukliansky, Tel Aviv University; Tova Milo, Tel Aviv University

Optimistic Recovery for Iterative Dataflows in Action

Sergey Doudoladov, Technische Universität Berlin; Chen Xu, Technische Universität Berlin; Sebastian Schelter, Technische Universität Berlin; Asterios Katsifodimos, Technische Universität Berlin; Stephan Ewen, Data Artisans GmbH; Kostas Tzoumas, Data Artisans GmbH; Volker Markl, Technische Universität Berlin

A Secure Search Engine for the Personal Cloud

Saliha Lallali, INRIA Rocquencourt; Nicolas Ancaux, INRIA Rocquencourt; Iulian Sandu Popa, INRIA Rocquencourt; Philippe Pucheral, INRIA Rocquencourt

A SQL Debugger Built from Spare Parts - Turning a SQL:1999 Database System into Its Own Debugger

Benjamin Dietrich, Universität Tübingen; Torsten Grust, Universität Tübingen

Just can't get enough - Synthesizing Big Data

Tilmann Rabl, University of Toronto; Manuel Danisch, bankmark; Michael Frank, bankmark; Sebastian Schindler, bankmark; Hans-Arno Jacobsen, University of Toronto

SIGMOD Tutorial 3.2: Data Management in Non-Volatile Memory (cont.)

Location: Plenary 1

Hashtag: #sigmod15 #tutorial3

Presenters: Stratis D. Viglas, University of Edinburgh

THURSDAY, 15:40 – 17:20

SIGMOD Research 23: Advanced Query Processing

Location: MR103

Hashtag: #sigmod15 #R23Px

Session Chair: Sharad Mehrotra (University of California, Irvine)

An Incremental Anytime Algorithm for Multi-Objective Query Optimization

Immanuel Trummer, EPFL; Christoph Koch, EPFL

Output-sensitive Evaluation of Prioritized Skyline Queries

Niccolo Meneghetti, University at Buffalo State University of New York; Denis Mindolin, Bloomberg L.P.; Paolo Ciaccia, University of Bologna; Jan Chomicki, University at Buffalo State University of New York

Learning Generalized Linear Models Over Normalized Data

Arun Kumar, University of Wisconsin-Madison; Jeffrey Naughton, University of Wisconsin-Madison; Jignesh Patel, University of Wisconsin-Madison

Utilizing IDs to Accelerate Incremental View Maintenance

Yannis Katsis, UC San Diego; Kian Win Ong, UC San Diego; Yannis Papakonstantinou, UC San Diego; Kevin Keliang Zhao, UC San Diego

SIGMOD Research 24: New Models

Location: MR104

Hashtag: #sigmod15 #R24Px

Session Chair: Jiaheng Lu (Renmin University of China)

S4: Top-k Spreadsheet-Style Search for Query Discovery

Fotis Psallidas, Columbia University; Bolin Ding, Microsoft; Kaushik Chakrabarti, Microsoft; Surajit Chaudhuri, Microsoft

Proactive Annotation Management in Relational Databases

Karim Ibrahim, Worcester Polytechnic Institute; Xiao Du, Worcester Polytechnic Institute; Mohamed Eltabakh, Worcester Polytechnic Institute

Weighted Coverage based Reviewer Assignment

Ngai Meng Kou, University of Macau; Leong Hou U, University of Macau; Nikos Mamoulis, University of Hong Kong; Zhiguo Gong, University of Macau

Distributed Online Tracking

Mingwang Tang, University of Utah; Feifei Li, University of Utah; Yufei Tao, Chinese University of Hong Kong

SIGMOD Demonstrations, Group A

Location: MR106

Hashtag: #sigmod15 #demoAx

CE-Storm: Confidential Elastic Processing of Data Streams

Nick Katsipoulakis, University of Pittsburgh; Cory Thoma, University of Pittsburgh; Eric Gratta, University of Pittsburgh; Alexandros Labrinidis, University of Pittsburgh; Adam Lee, University of Pittsburgh; Panos Chrysanthis, University of Pittsburgh

IReS: Intelligent, Multi-Engine Resource Scheduler for Big Data Analytics Workflows

Katerina Doka, National Technical University of Athens; Nikolaos Papailiou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Christos Mantas, National Technical University of Athens; Nectarios Koziris, National Technical University of Athens

Exploratory Keyword Search with Interactive Input

Zhifeng Bao, RMIT University; Yong Zeng, National University of Singapore; H. V. Jagadish, School of Computing; Tok Wang Ling, National University of Singapore

QE3D: Interactive Visualization and Exploration of Complex, Distributed Query Plans

Daniel Scheibll, SAP SE; Christian Dinse, SAP SE; Alexander Boehm, SAP SE

DataXFormer: An Interactive Data Transformation Tool

John Morcos, University of Waterloo; Ziawasch Abedjan, CSAIL MIT; Ihab Francis Ilyas, University of Waterloo; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Michael Stonebraker, MIT CSAIL

Quality-Driven Continuous Query Execution over Out-of-Order Data Streams

Yuanzhen Ji, SAP SE; Hongjin Zhou, SAP SE; Zbigniew Jerzak, SAP SE; Ioanaora Nica, SAP SE; Gregor Hackenbroich, SAP SE; Christof Fetzer, Technische Universitaet Dresden

MoDisSENSE: A Distributed Spatio-Temporal and Textual Processing Platform for Social Networking Services

Ioannis Mytilinis, National Technical University of Athens; Ioannis Giannakopoulos, National Technical University of Athens; Ioannis Konstantinou, National Technical University of Athens; Katerina Doka, National Technical University of Athens; Dimitrios Tsitsigkos, IMIS, RC Athena; Manolis Terrovitis, IMIS, RC Athena; Lampros Giampouras, Athens Technology Center S.A.; Nectarios Koziris, National Technical University of Athens

DocRicher: An Automatic Annotation System for Text Documents Using Social Media

Qiang Hu, SeSaMe Centre; Qi Liu, School of Computing; Xiaoli Wang, SeSaMe Centre; Anthony Tung, School of Computing; Shubham Goyal, SeSaMe Centre; Jisong Yang, SeSaMe Centre

A Demonstration of Rubato DB: A Highly Scalable NewSQL Database System for OLTP and Big Data Applications

Li-Yan Yuan, University of Alberta; Lengdong Wu, University of Alberta; Jia-Huai You, Computing Science; Yan Chi, Shanghai Shifang Software Ltd

G-OLA: Generalized On-Line Aggregation for Interactive Analysis on Big Data

Kai Zeng, University of California, Berkeley; Sameer Agarwal, Databricks Inc.; Ankur Dave, University of California, Berkeley; Michael Armbrust, Databricks Inc.; Ion Stoica, University of California, Berkeley

SIGMOD Tutorial 4: Knowledge Curation and Knowledge Fusion: Challenges, Models, and Applications

Location: MR105

Hashtag: #sigmod15 #tutorial4

Presenters: Xin Luna Dong, Google Inc.; Divesh Srivastava, AT&T Labs-Research

AWARDS

Edgar F. Codd Innovations Award

Presented to **Laura Haas** for pioneering innovations in the integration of multiple, heterogeneous data sources.



Laura Haas is an IBM Fellow and Director of IBM Research's Accelerated Discovery Lab. She was the director of computer science at IBM Almaden Research Center from 2005-2011, and had worldwide responsibility for IBM Research's exploratory science program from 2009 through 2013. Previously, Dr. Haas was responsible for Information Integration Solutions (IIS) architecture in IBM's Software Group after leading the IIS development team through its first two years. She joined the development team in 2001 as manager of DB2 UDB Query Compiler development. Before that, Dr. Haas was a research staff member and manager at the Almaden lab for nearly twenty years. In IBM Research, she worked on and managed a number of exploratory projects in distributed database systems. Dr. Haas is best known for her work on the Starburst query processor (from which DB2 UDB was developed); on Garlic, a system which allowed federation of heterogeneous data sources; and on Clio, the first semi-automatic tool for heterogeneous schema mapping. Garlic technology,

married with DB2 UDB query processing, is the basis for the IBM InfoSphere Federation Server, while Clio capabilities are a core differentiator in IBM's InfoSphere Data Architect. Dr. Haas is an active member of the database community. She served as Vice President of the VLDB Endowment Board of Trustees from 2004-2009 and was vice chair of ACM SIGMOD from 1989-1997. Dr. Haas has received several IBM awards for Outstanding Technical Achievement and Outstanding Innovation, and an IBM Corporate Award for her work on federated database technology. In 2010 she was recognized with the Anita Borg Institute Technical Leadership Award. She is a member of the National Academy of Engineering and the IBM Academy of Technology, an ACM Fellow, and Vice Chair of the board of the Computing Research Association. Dr. Haas received her PhD from the University of Texas at Austin, and her bachelor degree from Harvard University.

SIGMOD Systems Award

For fundamental contributions to the design and implementation of modern, object-relational database systems as embodied by the Postgres DBMS.



Dr. Michael Stonebraker has been a pioneer of data base research and technology for more than a quarter of a century. He was the main architect of the INGRES relational DBMS, and the object-relational DBMS, POSTGRES. These prototypes were developed at the University of California at Berkeley where Stonebraker was a Professor of Computer Science for twenty five years. More recently at M.I.T. he was a co-architect of the Aurora/Borealis stream processing engine, the C-Store column-oriented DBMS, the H-Store transaction processing engine, the SciDB array DBMS, and the Data Tamer data curation system. Presently he serves as Chief Technology Officer of Paradigm4 and Tamr, Inc.

Professor Stonebraker was awarded the ACM System Software Award in 1992 for his work on INGRES. Additionally, he was awarded the first annual SIGMOD Innovation award in 1994, and was elected to the National Academy of Engineering in 1997. He was awarded the IEEE John Von Neumann award in 2005, and is presently an Adjunct Professor of Computer Science at M.I.T., where he is co-director of the Intel Science and Technology Center focused on big data.



Lawrence A. Rowe is a Professor Emeritus of Electrical Engineering and Computer Science at UC Berkeley. He received a BA in mathematics and a PhD in information and computer science from UC Irvine in 1970 and 1976, respectively. He was a faculty member at Berkeley from 1976-2003. From 2007-2014, he served as President/COO and later Chairman/CEO at FX Palo Alto Laboratory. His research interests are multimedia systems and applications, distributed collaboration, DB application development tools, and distributed systems.

He is an ACM Fellow, past chair (1998-2003) of ACM SIGMM, and recipient of the 2009 ACM SIGMM Technical Achievement Award. He was a co-recipient of the 2002 UC Innovation in IT award for the development of the Berkeley Lecture Webcasting System and received a distinguished alumni award from UC Irvine in 2007.

Rowe is an active angel investor and co-founder of several companies including the original Ingres Corporation.

SIGMOD Contributions Award

Presented to **Curtis Dyreson** for extensive and outstanding service to the database community as ACM TODS and ACM SIGMOD Information Director and ACM SIGMOD Anthology and DiSC Editor.




Curtis Dyreson is an Associate Professor in the Department of Computer Science at Utah State University. His research interests include temporal databases, XML databases, data cubes, and aspect-oriented programming. He received his Ph.D. from the University of Arizona under the supervision of Richard T. Snodgrass. Prior to Utah State University he was on the faculty at James Cook University (Australia), Bond University (Australia), and Washington State University (USA), and was a visiting professor at Aalborg University (Denmark). He was the SIGMOD Anthology editor from 2004-2007, the SIGMOD DiSC editor from 2007-2010, the SIGMOD Information Director from 2011-now, the TODS Information Director from 2001-now, and since 2004 has served on the SIGMOD Executive Committee. He also served as the general co-chair of SIGMOD 2014 (with Feifei Li) and the program co-chair of DASFAA 2014 (with Sourav Bhowmick and Christian Jensen).

SIGMOD Test-of-Time Award

Robust and Fast Similarity Search for Moving Object Trajectories

Lei Chen, Hong Kong University of Science and Technology; M. Tamer Özsu, University of Waterloo; Vincent Oria, New Jersey Inst. of Technology.

An important consideration in similarity-based retrieval of moving object trajectories is the definition of a distance function. The existing distance functions are usually sensitive to noise, shifts and scaling of data that commonly occur due to sensor failures, errors in detection techniques, disturbance signals, and different sampling rates. Cleaning data to eliminate these is not always possible. In this paper, we introduce a novel distance function, Edit Distance on Real sequence (EDR) which is robust against these data imperfections. Analysis and comparison of EDR with other popular distance functions, such as Euclidean distance, Dynamic Time Warping (DTW), Edit distance with Real Penalty (ERP), and Longest Common Subsequences (LCSS), indicate that EDR is more robust than Euclidean distance, DTW and ERP, and it is on average 50% more accurate than LCSS. We also develop three pruning techniques to improve the retrieval efficiency of EDR and show that these techniques can be combined effectively in a search, increasing the pruning power significantly. The experimental results confirm the superior efficiency of the combined methods.



Lei Chen received the BS degree in computer science and engineering from Tianjin University, Tianjin, China, in 1994, the MA degree from Asian Institute of Technology, Bangkok, Thailand, in 1997, and the PhD degree in computer science from the University of Waterloo, Canada, in 2005. He is currently an associate professor in the Department of Computer Science and Engineering, Hong Kong University of Science and Technology. His research interests include crowdsourcing, probabilistic and uncertain databases, trajectory data analysis and privacy- preserved data publishing. The spatial crowdsourcing system, Gmission, developed by his team won the excellent demonstration award in VLDB 2014. He also got the best paper awards in DASFAA 2009 and 2010. Currently, he serves as associate editor-in-chief for IEEE Transaction on Data and Knowledge Engineering and a Trustee Board Member of VLDB Endowment.



M. Tamer Özsu is Professor of Computer Science at the University of Waterloo. His current research focuses on large scale data distribution and management of unconventional data (e.g., graphs, RDF, XML, streams).

He is a Fellow of ACM and IEEE, an elected member of the Science Academy (Turkey), and a member of Sigma Xi and AAAS. He was awarded the ACM SIGMOD Contributions Award in 2008 and the Ohio State University College of Engineering Distinguished Alumnus Award in 2008.



Vincent Oria is an associate professor of computer science at the New Jersey Institute of Technology (NJIT) in the USA. His research interests include multimedia databases, spatio-temporal databases, and similarity search in high-dimensional spaces.

He has held visiting professor positions at various institutions including National Institute of Informatics (Tokyo, Japan), Telecom-ParisTech (Paris, France), Université de Paris- Dauphine (Paris, France), INRIA (Roquencourt, France), Conservatoire National des Arts et Métiers (Paris,

France), Chinese University of Hong Kong (Hong Kong, China) and Hubei University of Science and Art (Xiang Yang, China). He is the recipient of the 2014 Outstanding Achievement in Research Award for the NJIT College of Computing Sciences.

SIGMOD Jim Gray Doctoral Dissertation Award

Dissertation Award winner: Alexander Thomson for the Dissertation Entitled "*Deterministic Transaction Execution in Distributed Database Systems*"

Dissertation Award honorable mention 1: Marina Drosou for the Dissertation Entitled "*Relevance and Diversity-based Ranking in Network-Centric Information Management Systems*";

Dissertation Award honorable mention 2: Karthik Ramachandra for the Dissertation Entitled "*Holistic Optimization of Database Applications*".



Alexander Thomson is a software engineer at Google, where he works on storage systems.



Marina Drosou holds a PhD from the Computer Science & Engineering Department of the University of Ioannina and BSc and MSc degrees from the Computer Science Department of the University of Ioannina. Her research interests include personalization via ranking based on diversity and recommendation systems. In the past, she has worked as an Instructor at the Department of Informatics and Telecommunications of the Technological Educational Institute of Epirus, Greece, and was also a visiting trainee researcher at the School of Information Technology and Electrical Engineering of the University of Queensland, Australia. She is currently a software engineer at the Information Systems Division of the Hellenic Police Headquarters. She has received scholarships from the Bodossaki Foundation and the Greek state ("Heraclitus II"). She is the recipient of the VLDB 2013 Best Paper Award. Her work has been published in TODS, VLDBJ, TKDE, PVLDB, etc.



Karthik Ramachandra is currently a Research SDE at the Microsoft Jim Gray Systems Lab. He received his Ph.D. in Computer Science from IIT Bombay, India, advised by Prof. S. Sudarshan. Karthik's dissertation, titled "*Holistic Optimization of Database Applications*", focuses on improving performance of database applications using techniques that span the boundaries of a traditional program compiler and a database query optimizer. It demonstrates that a combined holistic analysis of the program and its embedded queries opens up many hidden opportunities to optimize database applications. The thesis presents novel holistic techniques that automatically transform programs and queries to perform asynchronous prefetching and batching of queries in applications in order to optimize data access. It also presents a technique for optimization of user-defined functions in SQL queries. The thesis investigates practical issues that are of organic relevance to current database applications, and addresses these problems in a conceptually principled and technically novel manner. The application of programming language technologies towards optimizing database execution strategies, synergistically bringing together largely disparate streams of research activity, forms the core of this thesis.

SIGMOD Best Paper Award

DBSCAN Revisited: Mis-Claim, Un-Fixability, and Approximation

Junhao Gan, Chinese University of Hong Kong; Yufei Tao, Chinese University of Hong Kong;

DBSCAN is a popular method for clustering multi-dimensional objects. Just as notable as the method's vast success is the research community's quest for its efficient computation. The original KDD'96 paper claimed an algorithm with $O(n \log n)$ running time, where n is the number of objects. Unfortunately, this is a mis-claim; and that algorithm actually requires $O(n^2)$ time. There has been a fix in 2D space, where a genuine $O(n \log n)$ -time algorithm has been found. Looking for a fix for dimensionality $d \leq 3$ is currently an important open problem. In this paper, we prove that for $d \leq 3$, the DBSCAN problem requires $\Omega(n^{4/3})$ time to solve, unless very significant breakthroughs---ones widely believed to be impossible---could be made in theoretical computer science. This (i) explains why the community's search for fixing the aforementioned mis-claim has been futile for $d \leq 3$, and (ii) indicates (sadly) that *all* DBSCAN algorithms must be intolerably slow even on

moderately large n in practice. Surprisingly, we show that the running time can be dramatically brought down to $O(n)$ in expectation *regardless of the dimensionality d* , as soon as slight inaccuracy in the clustering results is permitted. We formalize our findings into the new notion of ρ -approximate DBSCAN, which we believe should replace DBSCAN on big data due to the latter's computational intractability.



Junhao Gan is a Ph. D. student in the Department of Computer Science and Engineering, The Chinese University of Hong Kong, supervised by Prof. Yupei Tao. He obtained his bachelor and master degrees from School of Software, Sun Yat-Sen University, in 2011 and 2013 respectively. His current interest aims to design practical algorithms with non-trivial theoretical guarantees for solving problems on massive data.



Yupei Tao is a full professor in the Department of Computer Science and Engineering, Chinese University of Hong Kong. He served as an associate editor of ACM Transactions on Database Systems (TODS) from 2008 to 2015, and of IEEE Transactions on Knowledge and Data Engineering (TKDE) from 2012 to 2014. He served as a PC co-chair of International Conference on Data Engineering (ICDE) 2014, and of International Symposium on Spatial and Temporal Databases (SSTD) 2011. His current research aims to leverage the theory of data structures and computational geometry to develop practical algorithms with non-trivial theoretical guarantees, particularly in dealing with massive datasets that do not fit in memory.

PODS Best Paper Award

Parallel-Correctness and Transferability for Conjunctive Queries

Tom J. Ameloot, Hasselt University; Gaetano Geck, TU Dortmund University; Bas Ketsman, Hasselt University; Frank Neven, Hasselt University; Thomas Schwentick, TU Dortmund University.

A dominant cost for query evaluation in modern massively distributed systems is the number of communication rounds. For this reason, there is a growing interest in single-round multiway join algorithms where data is first reshuffled over many servers and then evaluated in a parallel but communication-free way. The reshuffling itself is specified as a distribution policy. We introduce a correctness condition, called parallel-correctness, for the evaluation of queries w.r.t. a distribution policy. We study the complexity of parallel-correctness for conjunctive queries as well as transferability of parallel-correctness between queries. We also investigate the complexity of transferability for certain families of distribution policies, including, for instance, the Hypercube distribution.



Tom J. Ameloot is currently a Postdoctoral Fellow in the Databases and Theoretical Computer Science research group at Hasselt University (Belgium). Tom obtained his PhD on the topic of declarative networking under the supervision of Jan Van den Bussche. In his research, Tom likes to make theoretical insights for various research topics.



Frank Neven is a full professor at Hasselt University (Belgium). His main research area is that of database theory and systems. Recent research interests include big data, cloud computing, datamining in bioinformatics, automatic schema inference, and logic in databases.



Bas Ketsman is a Ph.D. student at Hasselt University under the supervision of Prof. Frank Neven. He is a PhD Fellow of the Research Foundation Flanders (FWO) and a member of the research group Databases and Theoretical Computer Science. His research interests include database theory and big data computations.



Thomas Schwentick works in several areas where Mathematical Logic meets Computer Science, very often on questions that concern the expressiveness of logical languages together the computational complexity of related algorithmic problems. He is Professor of Theoretical Computer Science at TU Dortmund University (Germany). He graduated in Mathematics the Johannes Gutenberg University in

Mainz, where he also finished his PhD and his habilitation. Before he moved to his current positions he had Professor positions in Jena and Marburg.



Gaetano Geck is a PhD student at TU Dortmund University, where he received his master's degree in computer science. He is interested in foundational aspects of Big Data and distributed databases.

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

XPath Satisfiability in the Presence of DTDs

Michael Benedikt, University of Oxford; Wenfei Fan, University of Edinburgh; Floris Geerts, University of Antwerp

We study the satisfiability problem associated with XPath in the presence of DTDs. This is the problem of determining, given a query p in an XPath fragment and a DTD D , whether or not there exists an XML document T such that T conforms to D and the answer of p on T is nonempty. We consider a variety of XPath fragments widely used in practice, and investigate the impact of different XPath operators on satisfiability analysis. We first study the problem for negation-free XPath fragments with and without upward axes, recursion and data-value joins, identifying which factors lead to tractability and which to NP-completeness. We then turn to fragments with negation but without data values, establishing lower and upper bounds in the absence and in the presence of upward modalities and recursion. We show that with negation the complexity ranges from PSPACE to EXPTIME. Moreover, when both data values and negation are in place, we find that the complexity ranges from NEXPTIME to undecidable. Finally, we give a finer analysis of the problem for particular classes of DTDs, exploring the impact of various DTD constructs, identifying tractable cases, as well as providing the complexity in the query size alone.



Michael Benedikt is Professor of Computer Science at Oxford University and a fellow of University College Oxford. He came to Oxford after a decade in US industrial research laboratories, including positions as Distinguished Member of Technical Staff at Bell Laboratories and visiting researcher at Yahoo! Labs. He has worked extensively in mathematical logic, finite model theory, verification, database theory, and database systems, and has served as chair of the ACM's main database theory conference, Principles of Database Systems. The current focus of his research is Web data management, with recent projects including querying of the deep Web, querying and integration of annotated data, and querying of web services.



Wenfei Fan is the Chair Professor of Web Data Management in the School of Informatics, University of Edinburgh, UK, and the director of the International Research Center on Big Data (RCBD), Beihang University, China. He is a Fellow of the Royal Society of Edinburgh, UK, a Fellow of the ACM, a National Professor of the 1000-Talent Program and a Yangtze River Scholar, China. He received his PhD from the University of Pennsylvania, and his MS and BS from Peking University. He is a recipient of the ACM PODS Alberto O. Mendelzon Test-of-Time Award in 2010, the Roger Needham Award in 2008 (UK), the NSF Career Award in 2001 (USA), and several Best Paper Awards (VLDB 2010, ICDE 2007, and Computer Networks 2002). His current research interests include database theory and systems, in particular big data, data quality, data integration, distributed query processing, recommender systems, social networks and Web services.



Floris Geerts holds a research professor position in the Advanced Database Research and Modelling group at the University of Antwerp, Belgium. Before that, he held a senior research fellow position in the database group at the University of Edinburgh, and a postdoc position in the data mining group at the University of Helsinki. He received his Ph.D in 2001 at the University of Hasselt, Belgium. His research interests include the theory and practice of databases and the study of data quality in particular. He is the (co-)author of 75+ articles that are mainly published in the proceedings of prestigious conferences (PODS, ICDT, VLDB, ICDE) and top-tier journals (ACM TODS, JCSS, JACM, TKDE). Many of his journal papers appeared in special issues that collect the best papers from databases conferences. In addition, Floris received two best paper awards (ICDM 2001, ICDE 2007). He is currently associate editor of ACM Transactions on Database Systems and was the general chair of EDBT/ICDT 2015.

Views and Queries: Determinacy and Rewriting

Luc Segoufin, INRIA & LSV Cachan; Victor Vianu, University of California San Diego

We investigate the question of whether a query Q can be answered using a set V of views. We first define the problem in information-theoretic terms: we say that V determines Q if V provides enough information to uniquely determine the answer to Q . Next, we look at the problem of rewriting Q in terms of V using a specific language. Given a view language V and query language Q , we say that a rewriting language R is complete for V -to- Q rewritings if every $Q \in \mathcal{Q}$ can be rewritten in terms of $V \in \mathcal{V}$ using a query in R , whenever V determines Q . While query rewriting using views has been extensively investigated for some specific languages, the connection to the information-theoretic notion of determinacy, and the question of completeness of a rewriting language, have received little attention. In this paper we investigate systematically the notion of determinacy and its connection to rewriting. The results concern decidability of determinacy for various view and query languages, as well as the power required of complete rewriting languages. We consider languages ranging from first-order to conjunctive queries.



Luc Segoufin got his PhD in 1999 from Université of Paris-Orsay. He was hired by INRIA since then. He is currently doing his research at the computer science department of ENS Cachan.



Victor Vianu received his PhD in Computer Science from the University of Southern California in 1983 and joined UC San Diego in 1984. Aside from UCSD, he has taught at the Ecole Normale Supérieure and Ecole Nationale Supérieure des Telecommunications in Paris, as well as the Sorbonne, and has spent several sabbaticals as invited professor at INRIA. Vianu's interests include database theory, computational logic, and Web data. His most recent research focuses on specification and verification of data-driven Web services and workflows. Vianu's publications include over 100 research articles and a graduate textbook on database theory. He received the PODS Alberto Mendelzon Test-of-Time Award in 2010 and has given numerous invited talks including keynotes at PODS, ICDT, STACS, the Association of Symbolic Logic, and the Federated Logic Conference. Vianu has served as General Chair of SIGMOD and Program Chair of the PODS and ICDT conferences. He is currently Editor-in-Chief of the Journal of the ACM and area editor of ACM Transactions on Computational Logic. He was elected Fellow of the ACM in 2006, Fellow of the AAAS in 2013, and is a member of Academia Europaea.

KEYNOTES

PODS Keynote: Computational Thinking, Inferential Thinking and "Big Data"

Michael I. Jordan, University of California, Berkeley

The phenomenon of "Big Data" is creating a need for research perspectives that blend computational thinking (with its focus on, e.g., abstractions, algorithms and scalability) with inferential thinking (with its focus on, e.g., underlying populations, sampling patterns, error bars and predictions). Database researchers and statistical machine learning researchers are centrally involved in the creation of this blend, and research that incorporates perspectives from both databases and machine learning will be of particular value in the bigger picture. This is true both for methodology and for theory. I present highlights of several research initiatives that draw jointly on database and statistical foundations, including work on concurrency control and distributed inference, subsampling, time/data tradeoffs and inference/privacy tradeoffs.



Michael I. Jordan is the Pehong Chen Distinguished Professor in the Department of Electrical Engineering and Computer Science and the Department of Statistics at the University of California, Berkeley. He received his Masters in Mathematics from Arizona State University, and earned his PhD in Cognitive Science in 1985 from the University of California, San Diego. He was a professor at MIT from 1988 to 1998. His research interests bridge the computational, statistical, cognitive and biological sciences. Prof. Jordan is a member of the National Academy of Sciences, a member of the National Academy of Engineering and a member of the American Academy of Arts and Sciences. He is a Fellow of the American Association for the Advancement of Science. He has been named a Neyman Lecturer and a Medallion Lecturer by the Institute of Mathematical Statistics. He received the David E. Rumelhart Prize in 2015 and the ACM/AAAI Allen Newell Award in 2009. He is a Fellow of the AAAI, ACM, ASA, CSS, IEEE, IMS, ISBA and SIAM.

SIGMOD Keynote Talk 1: From Data to Insights @ Bare Metal Speed

Jignesh M. Patel, University of Wisconsin-Madison

Data analytics platforms today largely employ data processing kernels (e.g. implementation of selection and join operator algorithms) that were developed for a now bygone hardware era. Hardware has made fundamental shifts in recent years, driven by the need to consider energy as a first-class design parameter. Consequently, across the processor-I/O hierarchy, the hardware paradigm today looks very different than it did just a few years ago. I argue that because of this shift, we are now building a "deficit" between the pace at which the hardware is evolving and the pace that is demanded of data processing kernels to keep up with the growth of big data. This deficit is unsustainable in the long run. One way to "pay off" this deficit is to have hardware and software co-evolve to exploit the full potential of the hardware. I will provide some examples of recent work from our Wisconsin Quickstep project that demonstrates the merit of this line of thinking. I'll focus on analytical data processing environments, and argue that our new way of storing data, called BitWeaving, and flattening databases into BitWeaved de-normalized tables, which we call WideTables, provides a dramatic new way to build "sustainable" analytical data processing systems. I will also discuss the implications of our approach on future hardware-software co-design for data analytics platforms.



Jignesh Patel is a Professor in Computer Sciences at the University of Wisconsin-Madison, where he also earned his Ph.D. He is the recipient of the Wisconsin COW teaching award, and the U. Michigan College of Engineering Education Excellence Award. He also has a strong interest in seeing research ideas transition to actual products. His Ph.D. thesis work was acquired by NCR/Teradata in 1997, and he also co-founded Locomatix -- a startup that built a platform to power real-time data-driven mobile services. Locomatix became part of Twitter in 2013. He is an ACM Fellow, serves on the board of Lands' End, and advises a number of startups. He blogs at <http://bigfastdata.blogspot.com>.

SIGMOD Keynote Talk 2: The Power Behind the Throne: Information Integration in the Age of Data-Driven Discovery

Laura M. Haas, IBM Research Accelerated Discovery Lab

Integrating data has always been a challenge. The information management community has made great progress in tackling this challenge, both on the theory and the practice. But in the last ten years, the world has changed dramatically. New platforms, devices and applications have made huge volumes of heterogeneous data available at speeds never contemplated before, while the quality of the available data has if anything degraded. Unstructured and semi-structured formats and no-sql data stores undercut the old reliable tools of schema, forcing applications to deal with data at the instance level. Deep expertise in the data and domain, in the tools and systems for integration and analysis, in mathematics, computer science, and business are needed to discover insights from data, but rarely are all of these skills found in a single individual or even team. Meanwhile, the availability of all these data has raised expectations for rapid breakthroughs in many sciences, for quick solutions to business problems, and for ever more sophisticated applications that combine and analyze information to solve our daily needs.

These expectations raise the bar for integration technology, while opening the door for it to play a broader role. Integration has always been a key player in handling data variety, for example, but now more than ever must deal with scale (in the number of types as well as in the volume and speed of data). While data cleansing has been one step of an integration pipeline, this technology must be leveraged throughout data integration, so that the integration process is better able to deal with the uncertainty in data, offering means to eliminate or reduce it, or, to elucidate it by linking important contextual information, such as provenance and usage. The complexity of today's data-driven challenges in fact suggests that the integration process should be context-aware, so that data sets may be combined differently depending on the proposed usage.

In the Accelerated Discovery Lab, we support data scientists working with a broad range of data as they try to find the insights to solve problems of business or societal importance. Clearly, integration is essential to insight. However, integration has to be across more than just datasets and schemas, and it has to be done more dynamically and flexibly than the standard tools allow. It is needed at multiple levels: (1) to build rich (but flexible) collections of diverse data, (2) to tightly bind individual data points into entities, allowing deeper explorations and (3) to bring together data and context to enable re-use by users with differing expertise. We think of the environment we are building as an integration hub for data, people and applications. It allows users to import, explore and create data and knowledge, inspired by the work of others, while it captures the patterns of decision-making and the provenance of decisions. I will describe the environment we are creating, the advances in the field that enable it, and the challenges that remain.



Laura Haas is an IBM Fellow and Director of IBM Research's Accelerated Discovery Lab. She was the director of computer science at IBM Almaden Research Center from 2005-2011, and had worldwide responsibility for IBM Research's exploratory science program from 2009 through 2013. Previously, Dr. Haas was responsible for Information Integration Solutions (IIS) architecture in IBM's Software Group after leading the IIS development team through its first two years. She joined the development team in 2001 as manager of DB2 UDB Query Compiler development. Before that, Dr. Haas was a research staff member and manager at the Almaden lab for nearly twenty years. In IBM Research, she worked on and managed a number of exploratory projects in distributed database systems. Dr. Haas is best known for her work on the Starburst query processor (from which DB2 UDB was developed); on Garlic, a system which allowed federation of heterogeneous data sources; and on Clio, the first semi-automatic tool for heterogeneous schema mapping. Garlic technology, married with DB2 UDB query processing, is the basis for the IBM InfoSphere Federation Server, while Clio capabilities are a core differentiator in IBM's InfoSphere Data Architect. Dr. Haas is an active member of the database community. She served as Vice President of the VLDB Endowment Board of Trustees from 2004-2009 and was vice chair of ACM SIGMOD from 1989-1997. Dr. Haas has received several IBM awards for Outstanding Technical Achievement and Outstanding Innovation, and an IBM Corporate Award for her work on federated database technology. In 2010 she was recognized with the Anita Borg Institute Technical Leadership Award. She is a member of the National Academy of Engineering and the IBM Academy of Technology, an ACM Fellow, and Vice Chair of the board of the Computing Research Association. Dr. Haas received her PhD from the University of Texas at Austin, and her bachelor degree from Harvard University.

TUTORIALS

PODS Tutorial 1: LogiQL, a Declarative Language for Enterprise Applications

Presenter: Todd J. Green, LogicBlox

We give an overview of LogiQL, a declarative, Datalog-based language for data management and analytics, along with techniques for efficient evaluation of LogiQL programs, emphasizing theoretical foundations when possible. These techniques include: leapfrog trijoin and its associated incremental maintenance algorithm, which we measure against appropriate optimality criteria; purely-functional data structures, which provide elegant versioning and branching capabilities that are indispensable for LogiQL; and transaction repair, a lock-free concurrency control scheme that uses LogiQL, incremental maintenance, and purely-functional data structures as essential ingredients.



T.J. Green is a Computer Scientist at LogicBlox. He received his Ph.D. in Computer and Information Science from the University of Pennsylvania in 2009. His awards include Best Student Paper at ICDT 2009, the Morris and Dorothy Rubinoff Award in 2010 (awarded to the outstanding computer science dissertation from the University of Pennsylvania), an honorable mention for the 2011 Jim Gray SIGMOD dissertation award, an NSF CAREER award in 2010, and Best Paper Runner-Up at ICDE 2012. Prior to LogicBlox, he was an Assistant Professor at UC Davis (2009-2011), an engineer at Xyleme (2001-2003), and a lead developer at Microsoft (1997-2001).

PODS Tutorial 2: Compact Summaries Over Large Datasets

Presenter: Graham Cormode, University of Warwick

It is often useful to build a compact summary data structure for a large data set, that will allow certain queries to be answered approximately. The database community has a rich history of building such summaries, in the form of histograms, wavelets and samples. However, it can sometimes be challenging to build such summaries when the data in question is large and distributed, and it is not convenient to sort or perform intensive processing of the data, even offline. A class of summaries for a variety of queries has emerged, drawing on algorithms that were originally created for streaming algorithms. These can be quite flexible: they can be computed for subsets of the data independently, then merged together to obtain a summary for the whole dataset. These allow approximate answering fundamental queries such as selectivity and frequency estimation, and newer summaries allow more complex functions from matrix and graph algorithms. This tutorial will introduce the concepts behind these summary methods, and outline several examples.



Graham Cormode is a Professor in Computer Science at the University of Warwick in the UK. He works on research topics in data management, privacy and big data analysis. Previously, he was a principal member of technical staff at AT&T Labs-Research from 2006-2013, and before this he was at Bell Labs, and Rutgers University. In 2013, he was recognized as a 'Distinguished Scientist' by the Association of Computing Machinery (ACM). His work has appeared in over 90 conference papers, 30 journal papers, and been awarded 25 US Patents. His work has received two best paper awards, and a ten year "Test of time" award for his work on sketching algorithms. He has edited two books on applications of algorithms to different areas, and coauthored a third. He currently serves as an associate editor for the ACM Transactions on Database Systems (TODS).

SIGMOD Tutorial 1: Overview of Data Exploration Techniques

Stratos Idreos, Harvard University; Olga Papaemmanouil, Brandeis University; Surajit Chaudhuri, Microsoft Research

Data exploration is about efficiently extracting knowledge from data even if we do not know exactly what we are looking for. In this tutorial, we survey recent developments in the emerging area of database systems tailored for data exploration. We discuss new ideas on how to store and access data as well as new ideas on how to interact with a data system to enable users and applications to quickly figure out which data parts are of interest. In addition, we discuss how to exploit lessons-learned from past research, the new challenges data exploration crafts, emerging applications and future research directions.



Stratos Idreos is an assistant professor of computer science at Harvard University where he leads DASlab, the Data Systems Laboratory at Harvard SEAS. Stratos works on data systems architectures with emphasis on designing systems for big data exploration that are easy to use and can stand the test of time. He received his PhD from University of Amsterdam while being a member of the CWI database group. For his doctoral work on Database Cracking, Stratos won the 2011 ACM SIGMOD Jim Gray Doctoral Dissertation award as well as the 2011 ERCIM Cor Baayen award as <most promising European young researcher in computer science and applied mathematics> from the European Research Council on Informatics and Mathematics. In addition, Stratos is a recipient of a 2010 IBM

Enterprise System Recognition award, a 2011 VLDB Challenges and Visions best paper award, a 2015 NSF CAREER award and the 2015 IEEE TCDE Early Career award.



Olga Papaemmanouil is an assistant professor of Computer Science at Brandeis University since 2009. She received her undergraduate degree from the University of Patras, Greece and completed her Ph.D at Brown University. Her research interests are in data management and distributed systems with a recent focus on performance management for cloud databases and interactive data exploration. She is the recipient of an NSF CAREER Award (2013) and a Paris Kanellakis Fellow (2002).



Surajit Chaudhuri is a Distinguished Scientist at Microsoft Research and leads the Data Management, Exploration and Mining group. In addition, as a Deputy Managing Director of Microsoft Research Lab at Redmond, he also has oversight of Distributed Systems, Networking, Security, Programming languages and Software Engineering groups of Microsoft Research, Redmond. He serves on the Senior Leadership Team of Microsoft's Cloud and Enterprises division. His current areas of interest are data discovery, self-manageability, and cloud database services. Working with his colleagues in Microsoft

Research, he helped incorporate the Index Tuning Wizard (and subsequently Database Engine Tuning Advisor) and data cleaning technology into Microsoft SQL Server. 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. Surajit received his Ph.D. from Stanford University in 1992.

SIGMOD Tutorial 2: Mining and Forecasting of Big Time-series Data

Presenters: Yasushi Sakurai, Kumamoto University; Yasuko Matsubara, Kumamoto University; Christos Faloutsos, Carnegie Mellon University

Given a large collection of time series, such as web-click logs, electric medical records and motion capture sensors, how can we efficiently and effectively find typical patterns? How can we statistically summarize all the sequences, and achieve a meaningful segmentation? What are the major tools for forecasting and outlier detection? Time-series data analysis is becoming of increasingly high importance, thanks to the decreasing cost of hardware and the increasing on-line processing capability.

The objective of this tutorial is to provide a concise and intuitive overview of the most important tools that can help us find patterns in large-scale time-series sequences. We review the state of the art in four related fields: (1) similarity search and pattern discovery, (2) linear modeling and summarization, (3) non-linear modeling and forecasting, and (4) the extension of time-series mining and tensor analysis. The emphasis of the tutorial is to provide the intuition behind these powerful tools, which is usually lost in the technical literature, as well as to introduce case studies that illustrate their practical use.



Yasushi Sakurai is a Professor at Kumamoto University. He obtained his B.E. degree from Doshisha University in 1991, and his M.E. and Ph.D. degrees from Nara Institute of Science and Technology in 1996 and 1999, respectively. In 1998, he joined NTT Laboratories, and became a Senior Research Scientist in 2005. He was a Visiting Researcher at Carnegie Mellon University during 2004-2005. He received two KDD best paper awards in 2008 and 2010. His research interests include time-series analysis, web mining, and sensor data processing.



Yasuko Matsubara is an Assistant Professor in the Department of Computer Science and Electrical Engineering at Kumamoto University, Japan. She obtained her BS and MS degrees from Ochanomizu University in 2007 and 2009 respectively, and her PhD from Kyoto University in 2012. She was a Visiting Researcher at Carnegie Mellon University during 2011-2012 and 2013-2014. Her research interests include time-series data mining and non-linear dynamic systems.



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), twenty "best paper" awards (including two "test of time" awards), and four teaching awards. Five of his advisees have attracted KDD or SCS dissertation awards. He is an ACM Fellow, he has served as a member of the executive committee of SIGKDD; he has published over 300 refereed articles, 17 book chapters and two monographs. He holds eight patents and he has given over 35 tutorials and over 15 invited distinguished lectures. His research interests include data mining for graphs and streams, fractals, database performance, and indexing for multimedia and bio-informatics data.

SIGMOD Tutorial 3: Data Management in Non-Volatile Memory **Presenters: Stratis D. Viglas, University of Edinburgh**

Non-volatile memory promises to bridge the gap between main memory and secondary storage by offering a universal storage device. Its performance profile is unique in that its latency is close to main memory and it is byte addressable, but it exhibits asymmetric I/O in that writes are more expensive than reads. These properties imply that it cannot act as a drop-in replacement for either main-memory or disk. Therefore, we must revisit the salient aspects of data management in light of this new technology. In what follows we present the current work in the area with a view towards identifying the open problems and exposing the research opportunities. In particular, we address issues like: (a) incorporating non-volatile memory into the data management stack, (b) supporting transactions and ensuring persistence and recovery, and (c) query processing.



Stratis Viglas is a Reader in Database Systems in the School of Informatics at the University of Edinburgh, which he joined after receiving his PhD from the University of Wisconsin-Madison. He has made contributions to data stream processing, XML data management, query processing and optimization, and data management over flash memory. His current work involves integrating managed runtimes with database systems for main memory query processing through just-in-time compilation of SQL queries and incorporating technologies like heterogeneous multicore and non-volatile memory into the data processing stack.

SIGMOD Tutorial 4: Knowledge Curation and Knowledge Fusion: Challenges, Models and Applications

Presenters: Xin Luna Dong, Google Inc.; Divesh Srivastava, AT&T Labs-Research

Large-scale knowledge repositories are becoming increasingly important as a foundation for enabling a wide variety of complex applications. In turn, building high-quality knowledge repositories critically depends on the technologies of knowledge curation and knowledge fusion, which share many similar goals with data integration, while facing even more challenges in extracting knowledge from both structured and unstructured data, across a large variety of domains, and in multiple languages.

Our tutorial highlights the similarities and differences between knowledge management and data integration, and has two goals. First, we introduce the Database community to the techniques proposed for the problems of entity linkage and relation extraction by the Knowledge Management, Natural Language Processing, and Machine Learning communities. Second, we give a detailed survey of the work done by these communities in knowledge fusion, which is critical to discover and clean errors present in sources and the many mistakes made in the process of knowledge extraction from sources. Our tutorial is example driven and hopes to build bridges between the Database community and other disciplines to advance research in this important area.



***Xin Luna Dong** is a Senior Research Scientist at Google Inc. She is one of the major contributors for the Knowledge Vault project, and has led the Knowledge-based Trust project, the Solomon data fusion project, and the Semex personal data management project. She has co-authored the book *Big Data Integration*, published over 50 papers in top conferences and journals, given over 10 tutorials, and got the Best Demo award in Sigmod 2005. She is the PC co-chair for WAIM 2015, is an associate editor of TKDE, and has served as an area chair for Sigmod 2015, ICDE 2013, and CIKM 2011.*



***Divesh Srivastava** is the head of Database Research at AT&T Labs-Research. He is a fellow of the Association for Computing Machinery (ACM), on the board of trustees of the VLDB Endowment, the managing editor of the *Proceedings of the VLDB Endowment (PVLDB)*, and an associate editor of the *ACM Transactions on Database Systems*. He received his Ph.D. from the University of Wisconsin, Madison, and his Bachelor of Technology from the Indian Institute of Technology, Bombay, India. His research interests and publications span a variety of topics in data management.*

SIGMOD NEW RESEARCHER SYMPOSIUM

Session Chairs: Alexandra Meliou, University of Massachusetts Amherst; Julia Stoyanovich, Drexel University

Every year, the SIGMOD Conference includes an exciting symposium addressing various aspects of starting a career in the data management community. These symposia are geared towards graduate students and junior researchers, but they attract a very broad audience because they are both informative and entertaining. Each panelist gives a 10-minute presentation, based on their personal experiences and perspectives gained over the years. After the presentations, an informal discussion ensues, where questions are welcomed from the audience.

This year's panel will address the topic: "How to fail." The panel will discuss various aspects of overcoming and learning from failure. This topic is of particular interest to graduate students and junior researchers, but is relevant to the entire data management community. The issues discussed in the presentation may include, but are not limited to:

- Bad research ideas: How to recognize them, avoid them, and abandon them. Can you transform a bad research direction into a good one?
- Whom should you ask for advice and should you always follow it?
- How should you balance your time across different responsibilities? Which responsibilities should you say "no" to? Is it OK to drop tasks that you committed to?
- What is the proper level of multitasking? How many research directions should you pursue at a time?
- When should you give up on a student, an advisor, or a topic?

SIGMOD RESEARCH PAPER ABSTRACTS

SIGMOD Research 1: Cloud: Parallel Execution

Distributed Outlier Detection using Compressive Sensing

Ying Yan, Microsoft Research; Jiaying Zhang, Microsoft Research; Bojun Huang, Microsoft Research; Xuzhan Sun, Peking University; Jiaqi Mu, University of Illinois, Urbana-Champaign; Zheng Zhang, NYU Shanghai; Thomas Moscibroda, Microsoft Research

Computing outliers and related statistical aggregation functions from large-scale big data sources is a critical operation in many cloud computing scenarios, e.g. service quality assurance, fraud detection, or novelty discovery. Such problems commonly have to be solved in a distributed environment where each node only has a local slice of the entirety of the data. To process a query on the global data, each node must transmit its local slice of data or an aggregated subset thereof to a global aggregator node, which can then compute the desired statistical aggregation function. In this context, reducing the total communication cost is often critical to the overall efficiency. In this paper, we show both theoretically and empirically that these communication costs can be significantly reduced for common distributed computing problems if we take advantage of the fact that production-level big data usually exhibits a form of sparse structure. Specifically, we devise a new aggregation paradigm for outlier detection and related queries. The paradigm leverages compressive sensing for data sketching in combination with outlier detection techniques. We further propose an algorithm that works even for non-sparse data that concentrates around an unknown value. In both cases, we show that the communication cost is reduced to the logarithm of the global data size. We incorporate our approach into Hadoop and evaluate it on real web-scale production data (distributed click-data logs). Our approach reduces data shuffling IO by up to 99%, and end-to-end job duration by up to 40% on many actual production queries.

Locality-aware Partitioning in Parallel Database Systems

Erfan Zamanian, Brown University; Carsten Binnig, Baden-Wuerttemberg Cooperative State University; Abdallah Salama, Baden-Wuerttemberg Cooperative

Parallel database systems horizontally partition large amounts of structured data in order to provide parallel data processing capabilities for analytical workloads in shared-nothing clusters. One major challenge when horizontally partitioning large amounts of data is to reduce the network costs for a given workload and a database schema. A common technique to reduce the network costs in parallel database systems is to co-partition tables on their join key in order to avoid expensive remote join operations. However, existing partitioning schemes are limited in that respect since only subsets of tables in complex schemata sharing the same join key can be co-partitioned unless tables are fully replicated. In this paper we present a novel partitioning scheme called predicate-based reference partition (or PREF for short) that allows to co-partition sets of tables based on given join predicates. Moreover, based on PREF, we present two automatic partitioning design algorithms to maximize data-locality. One algorithm only needs the schema and data whereas the other algorithm additionally takes the workload as input. In our experiments we show that our automated design algorithms can partition database schemata of different complexity and thus help to effectively reduce the runtime of queries under a given workload when compared to existing partitioning approaches.

ByteSlice: Pushing the Envelop of Main Memory Data Processing with a New Storage Layout

Ziqiang Feng, The Hong Kong Polytechnic University; Eric Lo, The Hong Kong Polytechnic University; Ben Kao, The University of Hong Kong; Wenjian Xu, The Hong Kong Polytechnic University

Scan and lookup are two core operations in main memory column stores. A scan operation scans a column and returns a result bit vector that indicates which records satisfy a filter. Once a column scan is completed, the result bit vector is converted into a list of record numbers, which is then used to look up values from other columns of interest for a query. Recently there are several in-memory data layout proposals that aim to improve the performance of in-memory data processing. However, these solutions all stand at either end of a trade-off --- each is either good in lookup performance or good in scan performance, but not both. In this paper we present ByteSlice, a new main memory storage layout that supports both highly efficient scans and lookups. ByteSlice is a byte-level columnar layout that fully leverages SIMD data-parallelism. Micro-benchmark experiments show that ByteSlice achieves a data scan speed at less than 0.5 processor cycle per column value --- a new limit of main memory data scan, without sacrificing lookup performance. Our experiments on TPC-H data and real data show that ByteSlice offers significant performance improvement over all state-of-the-art approaches.

Implicit Parallelism through Deep Language Embedding

Alexander Alexandrov, TU Berlin; Andreas Kuntz, TU Berlin; Asterios Katsifodimos, TU Berlin; Felix Schüler, TU Berlin; Lauritz Thamsen, TU Berlin; Odej Kao, TU Berlin; Tobias Herb, TU Berlin; Volker Markl, TU Berlin

The appeal of MapReduce has spawned a family of systems that implement or extend it. In order to enable parallel collection processing with User-Defined Functions (UDFs), these systems expose extensions of the MapReduce programming model as library-based dataflow APIs that are tightly coupled to their underlying runtime engine. Expressing data analysis algorithms with complex data and control flow structure using such APIs reveals a number of limitations that impede programmer's productivity. In this paper we show that the design of data analysis languages and APIs from a runtime engine point of view bloats the APIs with low-level primitives and affects programmer's productivity. Instead, we argue that an approach based on deeply embedding the APIs in a host language can address the shortcomings of current data analysis languages. To demonstrate this, we propose a language for complex data analysis embedded in Scala, which (i) allows for declarative specification of dataflows and (ii) hides the notion of data-parallelism and distributed runtime behind a suitable intermediate representation. We describe a compiler pipeline that facilitates efficient data-parallel processing without imposing runtime engine-bound syntactic or semantic restrictions on the structure of the input programs. We present a series of experiments with two state-of-the-art systems that demonstrate the optimization potential of our approach.

From Theory to Practice: Efficient Join Query Evaluation in a Parallel Database System

Shumo Chu, University of Washington; Magdalena Balazinska, University of Washington; Dan Suciu, University of Washington

Big data analytics often requires processing complex queries using massive parallelism, where the main performance metrics is the communication cost incurred during data reshuffling. In this paper, we describe a system that can compute efficiently complex join queries, including queries with cyclic joins, on a massively parallel architecture. We build on two independent lines of work for multi-join query evaluation: a communication-optimal algorithm for distributed evaluation, and a worst-case optimal algorithm for sequential evaluation. We evaluate these algorithms together, then describe novel, practical optimizations for both algorithms.

SIGMOD Research 2: Matrix and Array Computations

sPCA: Scalable Principal Component Analysis for Big Data on Distributed Platforms

Tarek Elgamal, Qatar Computing Research Institute; Maysam Yabandeh, Twitter; Ashraf Aboulnaga, Qatar Computing Research Institute; Walid Fathi, NTG Clarity; Mohamed Hefeeda, Qatar Computing Research Institute

Web sites, social networks, sensors, and scientific experiments currently generate massive amounts of data. Owners of this data strive to obtain insights from it, often by applying machine learning algorithms. Many machine learning algorithms, however, do not scale well to cope with the ever increasing volumes of data. To address this problem, we identify several optimizations that are crucial for scaling various machine learning algorithms in distributed settings. We apply these optimizations to the popular Principal Component Analysis (PCA) algorithm. PCA is an important tool in many areas including image processing, data visualization, information retrieval, and dimensionality reduction. We refer to the proposed optimized PCA algorithm as scalable PCA, or sPCA. sPCA achieves scalability via employing efficient large matrix operations, effectively leveraging matrix sparsity, and minimizing intermediate data. We implement sPCA on the widely-used MapReduce platform and on the memory-based Spark platform. We compare sPCA against the closest PCA implementations, which are the ones in Mahout/ MapReduce and MLlib/Spark. Our experiments show that sPCA outperforms both Mahout-PCA and MLlib-PCA by wide margins in terms of accuracy, running time, and volume of intermediate data generated during the computation.

Exploiting Matrix Dependency for Efficient Distributed Matrix Computation

Lele Yu, Peking University; Yingxia Shao, Peking University; Bin Cui, Peking University

Distributed matrix computation is a popular approach for many large-scale data analysis and machine learning tasks. However existing distributed matrix computation systems generally incur heavy communication cost during the runtime, which degrades the overall performance. In this paper, we propose a novel matrix computation system, named DMac, which exploits the matrix dependencies in matrix programs for efficient matrix computation in the distributed environment. We decompose each matrix program into a sequence of operations, and reveal the matrix dependencies between operations in the program. We next design a dependency-oriented cost model to select an optimal execution strategy for each operation, and generate a communication efficient execution plan for the matrix computation program. To facilitate the matrix computation in distributed systems, we further divide the execution plan into multiple un-interleaved stages which can run in a distributed cluster with efficient local execution strategy on each worker. The DMac system has been implemented on a

popular general-purpose data processing framework, Spark. The experimental results demonstrate that our techniques can significantly improve the performance of a wide range of matrix programs.

LEMP: Fast Retrieval of Large Entries in a Matrix Product

Christina Teflioudi, Max Planck Institute for Computer Science; Rainer Gemulla, University of Mannheim; Olga Mykytiuk, Sulzer GmbH

We study the problem of efficiently retrieving large entries in the product of two given matrices, which arises in a number of data mining and information retrieval tasks. We focus on the setting where the two input matrices are tall and skinny, i.e., with millions of rows and tens to hundreds of columns. In such settings, the product matrix is large and its complete computation is generally infeasible in practice. To address this problem, we propose the LEMP algorithm, which efficiently retrieves only the large entries in the product matrix without actually computing it. LEMP maps the large-entry retrieval problem to a set of smaller cosine similarity search problems, for which existing methods can be used. We also propose novel algorithms for cosine similarity search, which are tailored to our setting. Our experimental study on large real-world datasets indicates that LEMP is up to an order of magnitude faster than state-of-the-art approaches.

Skew-Aware Join Optimization for Array Databases

Jennie Duggan, Northwestern University; Olga Papaemmanouil, Brandeis University; Leilani Battle, MIT; Michael Stonebraker, MIT

Science applications are accumulating an ever-increasing amount of multidimensional data. Although some of it can be processed in a relational database, much of it is better suited to array-based engines. As such, it is important to optimize the query processing of these systems. This paper focuses on efficient query processing of join operations within an array database. These engines invariably “chunk” their data into multidimensional tiles that they use to efficiently process spatial queries. As such, traditional relational algorithms need to be substantially modified to take advantage of array tiles. Moreover, most n-dimensional science data is unevenly distributed in array space because its underlying observations rarely follow a uniform pattern. It is crucial that the optimization of array joins be skew-aware. In addition, owing to the scale of science applications, their query processing usually spans multiple nodes. This further complicates the planning of array joins. In this paper, we introduce a join optimization framework that is skew-aware for distributed joins. This optimization consists of two phases. In the first, a logical planner selects the query’s algorithm (e.g., merge join), the granularity of the its tiles, and the reorganization operations needed to align the data. The second phase implements this logical plan by assigning tiles to cluster nodes using an analytical cost model. Our experimental results, on both synthetic and real-world data, demonstrate that this optimization framework speeds up array joins by up to 2.5X in comparison to the baseline.

Resource Elasticity for Large-Scale Machine Learning

Botong Huang, Duke University; Matthias Boehm, IBM Research - Almaden; Yuanyuan Tian, IBM Research - Almaden; Berthold Reinwald, IBM Research - Almaden; Shirish Tatikonda, IBM Research - Almaden; Frederick Reiss, IBM Research - Almaden

Declarative large-scale machine learning (ML) aims at flexible specification of ML algorithms and automatic generation of hybrid runtime plans ranging from single node, in-memory computations to distributed computations on MapReduce (MR) or similar frameworks. State-of-the-art compilers in this context are very sensitive to memory constraints of the master process and MR cluster configuration. Different memory configurations can lead to significant performance differences. Interestingly, resource negotiation frameworks like YARN allow us to explicitly request preferred resources including memory. This capability enables automatic resource elasticity, which is not just important for performance but also removes the need for a static cluster configuration, which is always a compromise in multi-tenancy environments. In this paper, we introduce a simple and robust approach to automatic resource elasticity for large-scale ML. This includes (1) a resource optimizer to find near-optimal memory configurations for a given ML program, and (2) dynamic plan migration to adapt memory configurations during runtime. These techniques adapt resources according to data, program, and cluster characteristics. Our experiments demonstrate significant improvements up to 21x without unnecessary over-provisioning and low optimization overhead.

SIGMOD Research 3: Security and Access Control

SEMROD: Secure and Efficient MapReduce Over Hybrid Clouds

Kerim Yasin Oktay, University of California, Irvine; Vaibhav Khadilkar, The University of Texas at Dallas; Murat Kantarcioglu, The University of Texas at Dallas; Sharad Mehrotra, University of California, Irvine

This paper describes SEMROD, a sensitive data aware MapReduce (MR) framework for hybrid clouds. SEMROD steers data and computation through public and private machines in such a way that no knowledge about sensitive data is leaked to public machines. For this purpose, SEMROD keeps trace of intermediate keys (generated during MR execution) that become sensitive, based on which it makes dynamic task scheduling decisions. SEMROD guarantees that adversaries (*viz.*

public machines) cannot gain any “additional” information about sensitive data from either the data stored on public machines or the communication between public and private machines during job execution. SEMROD extends naturally from a single MR job to multi-phase MR jobs that result, for instance, from compiling Hive queries into MR jobs. Using SEMROD, computation that may involve sensitive data can exploit public machines, thereby bringing significant performance benefits. Such computation would otherwise be restricted to only private clouds. Our experiments clearly demonstrate performance advantages to using SEMROD as compared with other secure alternatives, even when the percentage of sensitive data is as high as 50%.

Authenticated Online Data Integration Services

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

Data integration involves combining data from multiple sources and providing users with a unified query interface. Data integrity has been a key problem in online data integration. Although a variety of techniques have been proposed to address the data consistency and reliability issues, there is little work on assuring the integrity of integrated data and the correctness of query results. In this paper, we take the first step to propose authenticated data integration services to ensure data and query integrity even in the presence of an untrusted integration server. We develop a novel authentication code called homomorphic secret sharing seal that can aggregate the inputs from individual sources faithfully by the untrusted server for future query authentication. Based on this, we design two authenticated index structures and authentication schemes for queries on multi-dimensional data. We further study the freshness problem in multi-source query authentication and propose several advanced update strategies. Analytical models and empirical results show that our seal design and authentication schemes are efficient and robust under various system settings.

ENKI: Access Control for Encrypted Query Processing

Isabelle Hang, University of Bremen; Florian Kerschbaum, SAP AG; Ernesto Damiani, University of Milan

A data owner outsourcing the database of a multi user application wants to prevent information leaks caused by outside attackers exploiting software vulnerabilities or by curious personnel. Query processing over encrypted data solves this problem for a single user, but provides only limited functionality in the face of access restrictions for multiple users and keys. ENKI is a system for securely executing queries over sensitive, access restricted data on an outsourced database. It introduces an encryption based access control model and techniques for query execution over encrypted, access restricted data on the database with only a few cases requiring computations on the client. A prototype of ENKI supports all queries seen in three real world use cases and executes queries from TPC-C benchmark with a modest overhead compared to the single user mode.

Collaborative Access Control in WebdamLog

Vera Zaychik Moffitt, Drexel University; Julia Stoyanovich, Drexel University; Serge Abiteboul, INRIA Saclay & ENS Cachan; Gerome Miklau, University of Massachusetts

The management of Web users' personal information is increasingly distributed across a broad array of applications and systems, including online social networks and cloud-based services. Users wish to share data using these systems, but avoiding the risks of unintended disclosures or unauthorized access by applications has become a major challenge. We propose a novel access control model that operates within a distributed data management framework based on datalog. Using this model, users can control access to data they own and control applications they run. They can conveniently specify access control policies providing flexible tuple-level control derived using provenance information. We present a formal specification of the model, an implementation built using an open-source distributed datalog engine, and an extensive experimental evaluation showing that the computational cost of access control is modest.

Automatic Enforcement of Data Use Policies with DataLawyer

Prasang Upadhyaya, University of Washington; Magdalena Balazinska, University of Washington; Dan Suciu, University of Washington

Data has value and is increasingly being exchanged for commercial and research purposes. Data, however, is typically accompanied by terms of use, which limit how it can be used. To date, there are only a few, ad-hoc methods to enforce these terms. We propose DataLawyer, a new system to formally specify usage policies and check them automatically at query runtime in a relational database management system (DBMS). We develop a new model to specify policies compactly and precisely. We introduce novel algorithms to efficiently evaluate policies that can cut policy-checking overheads to only a few percent of the total query runtime. We implement DataLawyer and evaluate it on a real database from the health-care domain.

SIGMOD Research 4: Cloud: Fault Tolerance, Reconfiguration

Cost-based Fault-tolerance for Parallel Data Processing

Abdallah Salama, Baden-Wuerttemberg Cooperative State University; Carsten Binnig, Baden-Wuerttemberg Cooperative State University; Tim Kraska, Brown University; Erfan Zamanian, Brown University

In order to deal with mid-query failures in parallel data engines (PDEs), different fault-tolerance schemes are implemented today: (1) fault-tolerance in parallel databases is typically implemented in a coarse-grained manner by restarting a query completely when a mid-query failure occurs, and (2) modern MapReduce-style PDEs implement a fine-grained fault-tolerance scheme, which either materializes intermediate results or implements a lineage model to recover from mid-query failures. However, neither of these schemes can efficiently handle mixed workloads with both short running interactive queries as well as long running batch queries nor do these schemes efficiently support a wide range of different cluster setups which vary in cluster size and other parameters such as the mean time between failures. In this paper, we present a novel cost-based fault-tolerance scheme which tackles this issue. Compared to the existing schemes, our scheme selects a subset of intermediates to be materialized such that the total query runtime is minimized under mid-query failures. Our experiments show that our cost-based fault-tolerance scheme outperforms all existing strategies and always selects the sweet spot for short- and long running queries as well as for different cluster setups.

Squall: Fine-Grained Live Reconfiguration for Partitioned Main Memory Databases

Aaron Elmore, University of Chicago; Vaibhav Arora, University of California, Santa Barbara; Rebecca Taft, MIT; Andrew Pavlo, Carnegie Mellon University; Divyakant Agrawal, University of California, Santa Barbara & Qatar Computing Research Institute; Amr El Abbadi, University of California, Santa Barbara

For data-intensive applications with many concurrent users, modern distributed main memory database management systems (DBMS) provide the necessary scale-out support beyond what is possible with single-node systems. These DBMSs are optimized for the short-lived transactions that are common in on-line transaction processing (OLTP) workloads. One way that they achieve this is to partition the database into disjoint subsets and use a single-threaded transaction manager per partition that executes transactions one-at-a-time in serial order. This minimizes the overhead of concurrency control mechanisms, but requires careful partitioning to limit distributed transactions that span multiple partitions. Previous methods used off-line analysis to determine how to partition data, but the dynamic nature of these applications means that they are prone to hotspots. In these situations, the DBMS needs to reconfigure how data is partitioned in real-time to maintain performance objectives. Bringing the system off-line to reorganize the database is unacceptable for on-line applications. To overcome this problem, we introduce the Squall technique for supporting live reconfiguration in partitioned, main memory DBMSs. Squall supports fine-grained repartitioning of databases in the presence of distributed transactions, high throughput client workloads, and replicated data. An evaluation of our approach on a distributed DBMS shows that Squall can reconfigure a database with no downtime and minimal overhead on transaction latency.

Madeus: Database Live Migration Middleware under Heavy Workloads for Cloud Environment

Takeshi Mishima, NTT; Yasuhiro Fujiwara, NTT

Database-as-a-service has been gaining popularity in cloud computing because multi-tenant databases can reduce costs by sharing off-the-shelf resources. However, due to heavy workloads, resource sharing often causes a hot spot; one node is overloaded even while others are not. Unfortunately, a hot spot can lead to violation of service level agreements and destroy customer satisfaction. To efficiently address the hot spot problem, we propose a middleware approach called Madeus that conducts database live migration. To make efficient database live migration possible, we also introduce the lazy snapshot isolation rule (LSIR) that enables concurrently propagating syncsets, which are the datasets needed to synchronize slave with master databases. Madeus provides efficient database live migration by implementing the LSIR under snapshot isolation. Unlike current approaches, Madeus is pure middleware that is transparent to the database management system and based on commodity hardware and software. To demonstrate the superiority of our approach over current approaches, we experimentally evaluated Madeus by using PostgreSQL with the TPC-W benchmark. The results indicate that Madeus achieves more efficient live migration than three other types of middleware approaches, especially under heavy workloads; therefore, it can effectively resolve hot spots.

Lineage-driven Fault Injection

Peter Alvaro, UC Berkeley; Joshua Rosen, UC Berkeley; Joseph Hellerstein, UC Berkeley

In large-scale data management systems, failure is practically a certainty. Fault-tolerant protocols and components are notoriously difficult to implement and debug. Worse still, choosing existing fault-tolerance mechanisms and integrating them correctly into complex systems remains an art form, and programmers have few tools to assist them. We propose a novel approach for discovering bugs in fault-tolerant data management systems: lineage-driven fault injection. A lineage-

driven fault injector reasons backwards from correct system outcomes to determine whether failures in the execution could have prevented the outcome. We present MOLLY, a prototype of lineage-driven fault injection that exploits a novel combination of data lineage techniques from the database literature and state-of-the-art satisfiability testing. If fault-tolerance bugs exist for a particular configuration, MOLLY finds them rapidly, in many cases using an order of magnitude fewer executions than random fault injection. Otherwise, MOLLY certifies that the code is bug-free for that configuration.

SIGMOD Research 5: Keyword Search and Text

Diversity-Aware Top-k Publish/Subscribe for Text Stream

Lisi Chen, Nanyang Technological University; Gao Cong, Nanyang Technological University

Massive amount of text data are being generated by a huge number of web users at an unprecedented scale. These data cover a wide range of topics. Users are interested in receiving a few up-to-date representative documents (e.g., tweets) that can provide them with a wide coverage of different aspects of their query topics. To address the problem, we consider the Diversity-Aware Top-k Subscription (DAS) query. Given a DAS query, we continuously maintain an up-to-date result set that contains k most recently returned documents over a text stream for the query. The DAS query takes into account text relevance, document recency, and result diversity. We propose a novel solution to efficiently processing a large number of DAS queries over a stream of documents. We demonstrate the efficiency of our approach on real-world dataset and the experimental results show that our solution is able to achieve a reduction of the processing time by 60–75% compared with two baselines. We also study the effectiveness of the DAS query.

Diverse and Proportional Size-1 Object Summaries for Keyword Search

Georgios Fakas, HKUST; Zhi Cai, Beijing University of Technology; Nikos Mamoulis, The University of Hong Kong

The abundance and ubiquity of graphs (e.g., Online Social Networks such as Google+ and Facebook; bibliographic graphs such as DBLP) necessitates the effective and efficient search over them. Given a set of keywords that can identify a Data Subject (DS), a recently proposed relational keyword search paradigm produces, as a query result, a set of Object Summaries (OSs). An OS is a tree structure rooted at the DS node (i.e., a tuple containing the keywords) with surrounding nodes that summarize all data held on the graph about the DS. OS snippets, denoted as size-1 OSs, have also been investigated. Size-1 OSs are partial OSs containing 1 nodes such that the summation of their importance scores results in the maximum possible total score. However, the set of nodes that maximize the total importance score may result in an uninformative size-1 OSs, as very important nodes may be repeated in it, dominating other representative information. In view of this limitation, in this paper we investigate the effective and efficient generation of two novel types of OS snippets, i.e. diverse and proportional size-1 OSs, denoted as DSize-1 and PSize-1 OSs. Namely, apart from the importance of each node, we also consider its frequency in the OS and its repetitions in the snippets. We conduct an extensive evaluation on two real graphs (DBLP and Google+). We verify effectiveness by collecting user feedback, e.g. by asking DBLP authors (i.e. the DSs themselves) to evaluate our results. In addition, we verify the efficiency of our algorithms and evaluate the quality of the snippets that they produce.

Local Filtering: Improving the Performance of Approximate Queries on String Collections

Xiaochun Yang, Northeastern University; Yaoshu Wang, University of New South Wales; Bin Wang, Northeastern University; Wei Wang, University of New South Wales

We study efficient query processing for approximate string queries, which find strings within a string collection whose edit distances to the query strings are within the given thresholds. Existing methods typically hinge on the property that globally *similar* strings must share at least certain number of *identical* substrings or subsequences. They become ineffective when there are burst errors or when the number of errors is large. In this paper, we explore the opposite paradigm focusing on finding out the *differences* of database strings to the query string. We propose a new filtering method, called *local filtering*, based on the idea that two strings exhibiting substantial local *dissimilarities* must be globally *dissimilar*. We propose the concept of (positional) local distance to quantify the minimum amount of errors a query fragment contributes to the edit distance between the query and a data string. It also leads to effective pruning rules and can speed up verification via early termination. We devise a family of indexing methods based on the idea of precomputing (positional) local distances for all possible combinations of query fragments and edit distance thresholds. Based on careful analyses of subtle relationships among local distances, novel techniques are proposed to drastically reduce the amount of enumeration with no or little impact on the pruning power. Efficient query processing methods exploiting the new index and bit-parallelism are also proposed. Experimental results on real datasets show that our local filtering-based methods can achieve substantial speedup compared with state-of-the-art methods, and they are robust against factors such as dataset characteristics and large edit distance thresholds.

Exact Top- k Nearest Keyword Search in Large Networks

Minhao Jiang, The Hong Kong University of Science and Technology; Ada Fu, The Chinese University of Hong Kong; Raymond Wong, The Hong Kong University of Science and Technology

Top- k nearest keyword search has been of interest because of applications ranging from road network location search by keyword to search of information on an RDF repository. We consider the evaluation of a query with a given vertex and a keyword, and the problem is to find a set of k nearest vertices that contain the keyword. The known algorithms for handling this problem only give approximate answers. In this paper, we propose algorithms for top- k nearest keyword search that provide exact solutions and which handle networks of very large sizes. We have also verified the performance of our solutions compared with the best-known approximation algorithms with experiments on real datasets.

Efficient Algorithms for Answering the m -Closest Keywords Query

Tao Guo, Nanyang Technological University; Xin Cao, Queen's University Belfast; Gao Cong, Nanyang Technological University

As an important type of spatial keyword query, the m -closest keywords (mCK) query finds a group of objects such that they cover all query keywords and have the smallest diameter, which is defined as the largest distance between any pair of objects in the group. The query is useful in many applications such as detecting locations of web resources. However, the existing work does not study the intractability of this problem and only provides exact algorithms, which are computationally expensive. In this paper, we prove that the problem of answering mCK queries is NP-hard. We first devise a greedy algorithm that has an approximation ratio of 2. Then, we observe that an mCK query can be approximately answered by finding the circle with the smallest diameter that encloses a group of objects together covering all query keywords. We prove that the group enclosed in the circle can answer the mCK query with an approximation ratio of $\frac{2}{\sqrt{3}}$. Based on this, we develop an algorithm for finding such a circle exactly, which has a high time complexity. To improve efficiency, we propose another two algorithms that find such a circle approximately, with a ratio of $(\frac{2}{\sqrt{3}} + \epsilon)$. Finally, we propose an exact algorithm that utilizes the group found by the $(\frac{2}{\sqrt{3}} + \epsilon)$ -approximation algorithm to obtain the optimal group. We conduct extensive experiments using real-life datasets. The experimental results offer insights into both efficiency and accuracy of the proposed approximation algorithms, and the results also demonstrate that our exact algorithm outperforms the best known algorithm by an order of magnitude.

SIGMOD Research 6: Graph Primitives

Minimum Spanning Trees in Temporal Graphs

SILU HUANG, Chinese University of Hong Kong; Ada Wai-Chee Fu, Chinese University of Hong Kong; Ruifeng Liu, Chinese University of Hong Kong

The computation of Minimum Spanning Trees (MSTs) is a fundamental graph problem with important applications. However, there has been little study of MSTs for temporal graphs, which is becoming common as time information is collected for many existing networks. We define two types of MSTs for temporal graphs, MST_a and MST_w , based on the optimization of time and cost, respectively. We propose efficient linear time algorithms for computing MST_a . We show that computing MST_w is much harder. We design efficient approximation algorithms based on a transformation to the Directed Steiner Tree problem (DST). Our solution also solves the classical DST problem with a better time complexity and the same approximation factor compared to the state-of-the-art algorithm. Our experiments on real temporal networks further verify the effectiveness of our algorithms. For MST_w , our solution is capable of shortening the runtime from 10 hours to 3 seconds.

Efficient Enumeration of Maximal k -Plexes

Devora Berlowitz, The Hebrew University of Jerusalem; Sara Cohen, The Hebrew University of Jerusalem; Benny Kimelfeld, Technion, Israel Institute of Technology

The problem of enumerating (i.e., generating) all maximal cliques in a graph has received extensive treatment, due to the plethora of applications in various areas such as data mining, bioinformatics, network analysis and community detection. However, requiring the enumerated subgraphs to be full cliques is too restrictive in common real-life scenarios where "almost cliques" are equally useful. Hence, the notion of a k -plex, a clique relaxation that allows every node to be "missing" k neighbors, has been introduced. But this seemingly minor relaxation casts existing algorithms for clique enumeration inapplicable, for inherent reasons. This paper presents the first provably efficient algorithms, both for enumerating the maximal k -plexes and for enumerating the maximal connected k -plexes. Our algorithms run in polynomial delay for a constant k and incremental FPT delay when k is a parameter. The importance of such algorithms is in the areas mentioned above, as well as in new applications. Extensive experimentation over both real and synthetic datasets shows the efficiency of our algorithms, and their scalability with respect to graph size, density and choice of k , as well as their clear superiority over the state-of-the-art.

Divide & Conquer: I/O Efficient Depth-First Search

Zhiwei Zhang, The Chinese University of Hong Kong; Jeffrey Xu Yu, The Chinese University of Hong Kong; Lu Qin, University of Technology, Sydney; Zechao Shang, The Chinese University of Hong Kong

Depth-First Search (DFS), which traverses a graph in the depth- first order, is one of the fundamental graph operations, and the result of DFS over all nodes in G is a spanning tree known as a DFS-Tree. There are many graph algorithms that need DFS such as connected component computation, topological sort, community detection, eulerian path computation, graph bipartiteness testing, planar graph testing, etc, because the in-memory DFS algorithm shows it can be done in linear time w.r.t. the size of G . However, given the fact that real-world graphs grow rapidly in the big data era, the in-memory DFS algorithm cannot be used to handle a large graph that cannot be entirely held in main memory. In this paper, we focus on I/O efficiency and study semi-external algorithms to DFS a graph G which is on disk. Here, like the existing semi-external algorithms, we assume that a spanning tree of G can be held in main memory and the remaining edges of G are kept on disk, and compute the DFS-Tree in main memory with which DFS can be identified. We propose novel divide & conquer algorithms to DFS over a graph G on disk. In brief, we divide a graph into several subgraphs, compute the DFS-Tree for each subgraph independently, and then merge them together to compute the DFS-Tree for the whole graph. With the global DFS-Tree computed we identify DFS. We discuss the valid division, that can lead to the correct DFS, and the challenges to do so. We propose two division algorithms, named Divide-Star and Divide-TD, and a merge algorithm. We conduct extensive experimental studies using four real massive datasets and several synthetic datasets to confirm the I/O efficiency of our approach.

Index-based Optimal Algorithms for Computing Steiner Components with Maximum Connectivity

Lijun Chang, University of New South Wales; Xuemin Lin, University of New South Wales; Lu Qin, University of Technology, Sydney; Jeffrey Xu Yu, The Chinese University of Hong Kong; Wenjie Zhang, University of New South Wales

With the proliferation of graph applications, the problem of efficiently computing all k -edge connected components of a graph G for a user-given k has been recently investigated. In this paper, we study the problem of efficiently computing the *steiner* component with the maximum connectivity; that is, given a set q of query vertices in a graph G , we aim to find the maximum induced subgraph g of G such that g contains q and g has the maximum connectivity, where g is denoted as smcc . To accommodate online query processing, we present an efficient algorithm based on a novel index such that the algorithm runs in linear time regarding the result size; thus, the algorithm is optimal since it needs at least linear time to output the result. Moreover, in this paper we also investigate variations of the above problem. We show that such a problem with the constraint that the size of the smcc is not smaller than a given size can also be solved in linear time regarding the result size (thus, optimal). We also show that the problem of computing the connectivity (rather than the graph details) of smcc can be solved in linear time regarding the query size (thus, optimal). To build the index, we extend the techniques in [\cite{SIGMOD13:Chang}](#) to accommodate batch processing and computation sharing. To efficiently support the applications with graph updates, we also present novel increment techniques. Finally, we conduct extensive performance studies on large real and synthetic graphs, which demonstrate that our index-based algorithms significantly outperform baseline algorithms by several orders of magnitude and our indexing algorithms are efficient.

SIGMOD Research 7: Data Mining

COMMIT: A Scalable Approach to Mining Communication Motifs from Dynamic Networks

Saket Gurukar, IIT Madras; Sayan Ranu, IIT Madras; Balaraman Ravindran, IIT Madras

A fundamental problem in behavioral analysis of human interactions is to understand how communications unfold. In this paper, we study this problem by mining Communication motifs from dynamic interaction networks. A communication motif is a recurring subgraph that has a similar sequence of information flow. Mining communication motifs requires us to explore the exponential subgraph search space where existing techniques fail to scale. To tackle this scalability bottleneck, we develop a technique called COMMIT. COMMIT converts a dynamic graph into a database of sequences. Through careful analysis in the sequence space, only a small portion of the exponential search space is accessed to identify regions embedding communication motifs. Extensive experiments on three different social networks show COMMIT to be up to two orders of magnitude faster than baseline techniques. Furthermore, qualitative analysis demonstrate communication motifs to be effective in characterizing the recurring patterns of interactions while also revealing the role that the underlying social network plays in shaping human behavior.

LASH: Large-Scale Sequence Mining with Hierarchies

Kaustubh Beedkar, University of Mannheim; Rainer Gemulla, University of Mannheim

We propose LASH, a scalable, distributed algorithm for mining sequential patterns in the presence of hierarchies. LASH takes as input a collection of sequences, each composed of items from some application-specific vocabulary. In contrast to traditional approaches to sequence mining, the items in the vocabulary are arranged in a hierarchy; both input sequences and sequential patterns may consist of items from different levels of the hierarchy. Such hierarchies naturally occur in a number of applications including mining natural-language text, customer transactions, error logs, or event sequences. LASH is the first parallel algorithm for mining frequent sequences with hierarchies; it is designed to scale to very large datasets. At its heart, LASH partitions the data using a novel, hierarchy-aware variant of item-based partitioning and subsequently mines each partition independently and in parallel using a customized mining algorithm called pivot sequence miner. LASH is amenable to a MapReduce implementation; we propose effective and efficient algorithms for both the construction and the actual mining of partitions. Our experimental study on large real-world datasets suggest good scalability and run-time efficiency.

Twister Tries: Approximate Hierarchical Agglomerative Clustering for Average Distance in Linear Time

Michael Cochez, University of Jyväskylä; Hao Mou, University of Jyväskylä

Many commonly used data-mining techniques utilized across research fields perform poorly when used for large data sets. Sequential agglomerative hierarchical non-overlapping clustering is one technique for which the algorithms' scaling properties prohibit clustering of a large amount of items. Besides the unfavorable time complexity of $O(n^2)$, these algorithms have a space complexity of $O(n^2)$, which can be reduced to $O(n)$ if the time complexity is allowed to rise to $O(n^2 \log^2 n)$. In this paper, we propose the use of locality-sensitive hashing combined with a novel data structure called twister tries to provide an approximate clustering for average linkage. Our approach requires only linear space. Furthermore, its time complexity is linear in the number of items to be clustered, making it feasible to apply it on a larger scale. We evaluate the approach both analytically and by applying it to several data sets.

DBSCAN Revisited: Mis-Claim, Un-Fixability, and Approximation

Junhao Gan, Chinese University of Hong Kong; Yufei Tao, Chinese University of Hong Kong

DBSCAN is a popular method for clustering multi-dimensional objects. Just as notable as the method's vast success is the research community's quest for its efficient computation. The original KDD'96 paper claimed an algorithm with $O(n \log n)$ running time, where n is the number of objects. Unfortunately, this is a mis-claim; and that algorithm actually requires $O(n^2)$ time. There has been a fix in 2D space, where a genuine $O(n \log n)$ -time algorithm has been found. Looking for a fix for dimensionality $d \leq 3$ is currently an important open problem. In this paper, we prove that for $d \leq 3$, the DBSCAN problem requires $\Omega(n^{d/2})$ time to solve, unless very significant breakthroughs---ones widely believed to be impossible---could be made in theoretical computer science. This (i) explains why the community's search for fixing the aforementioned mis-claim has been futile for $d \leq 3$, and (ii) indicates (sadly) that *all* DBSCAN algorithms must be intolerably slow even on moderately large n in practice. Surprisingly, we show that the running time can be dramatically brought down to $O(n)$ in expectation *regardless of the dimensionality d* , as soon as slight inaccuracy in the clustering results is permitted. We formalize our findings into the new notion of p -approximate DBSCAN, which we believe should replace DBSCAN on big data due to the latter's computational intractability.

The TagAdvisor: Luring the Lurkers to Review Web Items

Azade Nazi, University of Texas at Arlington; Mahashweta Das, HP Labs; Gautam Das, University of Texas at Arlington

The increasing popularity and widespread use of online review sites over the past decade has motivated businesses of all types to possess an expansive arsenal of user feedback (preferably positive) in order to mark their reputation and presence in the Web. Though a significant proportion of purchasing decisions today are driven by average numeric scores (e.g., movie rating in IMDB), detailed reviews are critical for activities such as buying an expensive digital SLR camera, reserving a vacation package, etc. Since writing a detailed review for a product (or, a service) is usually time-consuming and may not offer any incentive, the number of useful reviews available in the Web is far from many. The corpus of reviews available at our disposal for making informed decisions also suffers from spam and misleading content, typographical and grammatical errors, etc. In this paper, we address the problem of how to engage the lurkers (i.e., people who read reviews but never take time and effort to write one) to participate and write online reviews by systematically simplifying the reviewing task. Given a user and an item that she wants to review, the task is to identify the top- k meaningful phrases (i.e., tags) from the set of all tags (i.e., available user feedback for items) that, when advised, would help her review an item easily. We refer to it as the TagAdvisor problem, and formulate it as a general-constrained optimization goal. Our framework is centered around three measures - relevance (i.e., how well the result set of tags describes an item to a user), coverage (i.e., how well the result set of tags covers the different aspects of an item), and polarity (i.e., how well sentiment is attached to the result set of

tags) in order to help a user review an item satisfactorily. By adopting different definitions of coverage, we identify two concrete problem instances that enable a wide range of real-world scenarios. We show that these problems are NP-hard and develop practical algorithms with theoretical bounds to solve them efficiently. We conduct detailed experiments on synthetic and real data crawled from the web to validate the utility of our problem and effectiveness of our solutions.

SIGMOD Research 8: Uncertainty and Linking

Supporting Data Uncertainty in Array Databases

Liping Peng, University of Massachusetts Amherst; Yanlei Diao, University of Massachusetts Amherst

Uncertain data management has become crucial to scientific applications. Recently, array databases have gained popularity for scientific data processing due to performance benefits. In this paper, we address uncertain data management in array databases, which may involve both value uncertainty within individual tuples and position uncertainty regarding where a tuple should belong in an array given uncertain dimension attributes. Our work defines the formal semantics of array operations under both value and position uncertainty. To address the new challenge raised by position uncertainty, we propose a suite of storage and evaluation strategies for array operations, with a focus on a new scheme that bounds the overhead of querying by strategically treating tuples with large variances via replication in storage. Results from real datasets show that for common workloads, our best-performing techniques outperform alternative methods based on state-of-the-art indexes by 1.7x to 4.3x for the Subarray operation and 1 to 2 orders of magnitude for Structure-Join, at only a small storage cost.

Identifying the Extent of Completeness of Query Answers over Partially Complete Databases

Simon Razniewski, Free University of Bozen-Bolzano; Flip Korn, Google Research; Werner Nutt, Free University of Bozen-Bolzano; Divesh Srivastava, AT&T Labs-Research

In many applications including loosely coupled cloud databases, collaborative editing and network monitoring, data from multiple sources is regularly used for query answering. For reasons such as system failures, insufficient author knowledge or network issues, data may be temporarily unavailable or generally nonexistent. Hence, not all data needed for query answering may be available. In this paper, we propose a natural class of completeness patterns, expressed by selections on database tables, to specify complete parts of database tables. We then show how to adapt the operators of relational algebra so that they manipulate these completeness patterns to compute completeness patterns pertaining to query answers. Our proposed algebra is computationally sound and complete with respect to the information that the patterns provide. We show that stronger completeness patterns can be obtained by considering not only the schema but also the database instance and we extend the algebra to take into account this additional information. We develop novel techniques to efficiently implement the computation of completeness patterns on query answers and demonstrate their scalability on real data.

k-Hit Query: Top-k Query with Probabilistic Utility Function

Peng Peng, Hong Kong University of Science and Technology; Raymond Chi-Wing Wong, Hong Kong University of Science and Technology

Multi-criteria decision making problem has been well studied for many years. One popular query for multi-criteria decision making is top-k queries which require each user to specify an *exact* utility function. In many cases, the utility function of a user is *probabilistic* and finding the distribution on the utility functions has been widely explored in the machine learning areas, such as user's recommender systems, Bayesian learning models and user's preference elicitation, for improving user's experience. Motivated by this, we propose a new type of queries called *k-hit* queries, which has not been studied before. Given a set D of tuples in the database, the distribution θ on utility functions and a positive integer k , we would like to select a set of k tuples from D in order to maximize the probability that at least one of tuples in the selection set is the favorite of a user. All applications for top- k queries can naturally be used in *k-hit* queries. In this paper, we present various interesting properties of *k-hit* queries. Besides, based on these properties, we propose a novel algorithm called *k-hit_{Alg}* for *k-hit* queries. Finally, we conducted comprehensive experiments to show that the performance of our proposed method, *k-hit_{Alg}*, is superior compared with other existing algorithms which were originally used to answer other existing queries.

Linking Temporal Records for Profiling Entities

Furong Li, National University of Singapore; Mong Li Lee, National University of Singapore; Wynne Hsu, National University of Singapore; Wang-Chiew Tan, UC Santa Cruz

To harness the rich amount of information available on the Web today, many organizations start to aggregate public (and private) data to derive new knowledge bases. A fundamental challenge in constructing an accurate integrated knowledge repository from different data sources is to understand how facts across different sources are related to one another over time. This challenge, referred to as the temporal record linkage problem, goes far beyond the traditional record linkage

problem as it requires a fine-grained analysis of how two facts are temporally related if they both refer to the same entity. In this paper, we present a new solution for understanding how two facts may be temporally related and exploit the knowledge to profile how entities evolve over time. Our solution makes use of a novel transition model which captures sophisticated patterns of value transitions. Specifically, our transition model captures the probability that an entity may change to a particular attribute value after some time period. This transition model can be considered jointly with various source quality metrics to fine-tune how records should be temporally linked to entities. In particular, we showcase how the freshness of data sources can be built into a source-aware temporal matching algorithm that jointly considers the value transitions and the freshness of data sources to link temporal records to entities in the right time period. In this way, an increasingly complete and up-to-date entity profile can be derived as more and more temporal records are aggregated from different sources. Our suite of experimental results on real world datasets demonstrate that our proposed method is able to outperform the state-of-the-art techniques and build more complete profiles for entities by identifying their true matching temporal records at the right time period.

SIGMOD Research 9: Transactional Architectures

On the Design and Scalability of Distributed Shared-Data Databases

Simon Loesing, ETH Zurich; Markus Pilman, ETH Zurich; Thomas Etter, ETH Zurich; Donald Kossmann, Microsoft Research

Database scale-out is commonly implemented by partitioning data across several database instances. This approach, however, has several restrictions. In particular, partitioned databases are inflexible in large-scale deployments and assume a partition-friendly workload in order to scale. In this paper, we analyze an alternative architecture design for distributed relational databases that overcomes the limitations of partitioned databases. The architecture is based on two fundamental principles: We decouple query processing and transaction management from data storage, and we share data across query processing nodes. The combination of these design choices provides scalability, elasticity, and operational flexibility without making any assumptions on the workload. As a drawback, sharing data among multiple database nodes causes synchronization overhead. To address this limitation, we introduce techniques for scalable transaction processing in shared-data environments. Specifically, we describe mechanisms for efficient data access, concurrency control, and data buffering. In combination with new hardware trends, the techniques enable performance characteristics that top state-of-the-art partitioned databases.

Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems

Thomas Neumann, Technische Universität München; Tobias Mühlbauer, Technische Universität München; Alfons Kemper, Technische Universität München

Multi-Version Concurrency Control (MVCC) is a widely employed concurrency control mechanism, as it allows for execution modes where readers never block writers. However, most systems implement only snapshot isolation (SI) instead of full serializability. Adding serializability guarantees to existing SI implementations tends to be prohibitively expensive. We present a novel MVCC implementation for main-memory database systems that has very little overhead compared to serial execution with single-version concurrency control, even when maintaining serializability guarantees. Updating data in-place and storing versions as before-image deltas in undo buffers not only allows us to retain the high scan performance of single-version systems but also forms the basis of our cheap and fine-grained serializability validation mechanism. The novel idea is based on an adaptation of precision locking and verifies that the (extensional) writes of recently committed transactions do not intersect with the (intensional) read predicate space of a committing transaction. We experimentally show that our MVCC model allows very fast processing of transactions with point accesses as well as read-heavy transactions and that there is little need to prefer SI over full serializability any longer.

FOEDUS: OLTP Engine for a Thousand Cores and NVRAM

Hideaki Kimura, HP Labs

Server hardware is about to drastically change. As typified by emerging hardware such as UC Berkeley's Firebox project and by Intel's Rack-Scale Architecture (RSA), next generation servers will have *thousands of cores*, *large DRAM*, and *huge NVRAM*. We analyze the characteristics of these machines and find that no existing database is appropriate. Hence, we are developing *FOEDUS*, an open-source, from-scratch database engine whose architecture is drastically different from traditional databases. It extends in-memory database technologies to further scale up and also allows transactions to efficiently manipulate data pages in both DRAM and NVRAM. We evaluate the performance of FOEDUS in a large NUMA machine (16 sockets and 240 physical cores) and find that FOEDUS achieves multiple orders of magnitude higher TPC-C throughput compared to H-Store with anti-caching.

Let's Talk About Storage & Recovery Methods for Non-Volatile Memory Database Systems

Joy Arulraj, Carnegie Mellon University; Andrew Pavlo, Carnegie Mellon University; Subramanya Dullloor, Intel Labs

The advent of non-volatile memory (NVM) will fundamentally change the dichotomy between memory and durable storage in database management systems (DBMSs). These new NVM devices are almost as fast as DRAM, but all writes to it are potentially persistent even after power loss. Existing DBMSs are unable to take full advantage of this technology because their internal architectures are predicated on the assumption that memory is volatile. With NVM, many of the components of legacy DBMSs are unnecessary and will degrade the performance of data intensive applications. To better understand these issues, we implemented three engines in a modular DBMS testbed that are based on different storage management architectures: (1) in-place updates, (2) copy-on-write updates, and (3) log-structured updates. We then present NVM-aware variants of these architectures that leverage the persistence and byte-addressability properties of NVM in their storage and recovery methods. Our experimental evaluation on an NVM hardware emulator shows that these engines achieve up to 5.5X higher throughput than their traditional counterparts while reducing the amount of wear due to write operations by up to 2X. We also demonstrate that our NVM-aware recovery protocols allow these engines to recover almost instantaneously after the DBMS restarts.

SIGMOD Research 10: Privacy

Private Release of Graph Statistics using Ladder Functions

Jun Zhang, Nanyang Technological University; Graham Cormode, University of Warwick; Cecilia Procopiuc, Google Inc.; Divesh Srivastava, AT&T Labs - Research; Xiaokui Xiao, Nanyang Technological University

Protecting the privacy of individuals in graph structured data while making accurate versions of the data available is one of the most challenging problems in data privacy. Most efforts to date to perform this data release end up mired in complexity, overwhelm the signal with noise, and are not effective for use in practice. In this paper, we introduce a new method which guarantees differential privacy. It specifies a probability distribution over possible outputs that is carefully defined to maximize the utility for the given input, while still providing the required privacy level. The distribution is designed to form a 'ladder', so that each output achieves the highest 'rung' (maximum probability) compared to less preferable outputs. We show how our ladder framework can be applied to problems of counting the number of occurrences of subgraphs, a vital objective in graph analysis, and give algorithms whose cost is comparable to that of computing the count exactly. Our experimental study confirms that our method outperforms existing methods for counting triangles and stars in terms of accuracy, and provides solutions for some problems for which no effective method was previously known. The results of our algorithms can be used to estimate the parameters of suitable graph models, allowing synthetic graphs to be sampled.

Bayesian Differential Privacy on Correlated Data

Bin Yang, Rakuten Inc.; Issei Sato, The University of Tokyo; Hiroshi Nakagawa, The University of Tokyo

Differential privacy provides a rigorous standard for evaluating the privacy of perturbation algorithms. It has widely been regarded that differential privacy is a universal definition that deals with both independent and correlated data and a differentially private algorithm can protect privacy against arbitrary adversaries. However, recent research indicates that differential privacy may not guarantee privacy against arbitrary adversaries if the data are correlated. In this paper, we focus on the private perturbation algorithms on correlated data. We investigate the following three problems: (1) the influence of data correlations on privacy; (2) the influence of adversary prior knowledge on privacy; and (3) a general perturbation algorithm that is private for prior knowledge of any subset of tuples in the data when the data are correlated. We propose a Pufferfish definition of privacy, called Bayesian differential privacy, by which the privacy level of a probabilistic perturbation algorithm can be evaluated even when the data are correlated and when the prior knowledge is incomplete. We present a Gaussian correlation model to accurately describe the structure of data correlations and analyze the Bayesian differential privacy of the perturbation algorithm on the basis of this model. Our results show that privacy is poorest for an adversary who has the least prior knowledge. We further extend this model to a more general one that considers uncertain prior knowledge.

Modular Order-Preserving Encryption, Revisited

Charalampos Mavroforakis, Boston University; Nathan Chenette, Rose-Hulman Institute of Technology; Adam O'Neill, Georgetown University and NIST; George Kollios, Boston University; Ran Canetti, Boston University

Order-preserving encryption (OPE) schemes, whose ciphertexts preserve the natural ordering of the plaintexts, allow efficient range query processing over outsourced encrypted databases without giving the server access to the decryption key. Such schemes have recently received increased interest in both the database and the cryptographic communities. In

particular, modular order-preserving encryption (MOPE), due to Boldyreva et al., is a promising extension that increases the security of the basic OPE by introducing a secret modular offset to each data value prior to encrypting it. However, executing range queries via MOPE in a naive way allows the adversary to learn this offset, negating any potential security gains of this approach. In this paper, we systematically address this vulnerability and show that MOPE can be used to build a practical system for executing range queries on encrypted data while providing a significant security improvement over the basic OPE. We introduce two new query execution algorithms for MOPE: our first algorithm is efficient if the user's query distribution is well-spread, while the second scheme is efficient even for skewed query distributions. Interestingly, our second algorithm achieves this efficiency by leaking the least-important bits of the data, whereas OPE is known to leak the most-important bits of the data. We also show that our algorithms can be extended to the case where the query distribution is adaptively learned online. We present new, appropriate security models for MOPE and use them to rigorously analyze the security of our proposed schemes. Finally, we design a system prototype that integrates our schemes on top of an existing database system and apply query optimization methods to execute SQL queries with range predicates efficiently. We provide a performance evaluation of our prototype under a number of different database and query distributions, using both synthetic and real datasets.

Chiaroscuro: Transparency and Privacy for Massive Personal Time-Series Clustering

Tristan Allard, IRISA & Univ. Rennes 1; Georges Hébrail, EDF R&D; Florent Massegla, Inria & Lirmm, Univ. Montpellier; Esther Pacitti, Inria & Lirmm, Univ. Montpellier

The advent of on-body/at-home sensors connected to personal devices leads to the generation of fine grain highly sensitive personal data at an unprecedented rate. However, despite the promises of large scale analytics there are obvious privacy concerns that prevent individuals to share their personal data. In this paper, we propose Chiaroscuro, a complete solution for clustering personal data with strong privacy guarantees. The execution sequence produced by Chiaroscuro is massively distributed on personal devices, coping with arbitrary connections and disconnections. Chiaroscuro builds on our novel data structure, called Diptych, which allows the participating devices to collaborate privately by combining encryption with differential privacy. Our solution yields a high clustering quality while minimizing the impact of the differentially private perturbation. Chiaroscuro is both correct and secure. Finally, we provide an experimental validation of our approach on both real and synthetic sets of time-series.

SIGMOD Research 11: Streams

Persistent Data Sketching

Zhewei Wei, Renmin University of China; Ge Luo, The Hong Kong University of Science and Technology; Ke Yi, The Hong Kong University of Science and Technology; Xiaoyong Du, Renmin University of China; Ji-Rong Wen, Renmin University of China

A *persistent data structure*, also known as a *multiversion data structure* in the database literature, is a data structure that preserves all its previous versions as it is updated over time. Every update (inserting, deleting, or changing a data record) to the data structure creates a new version, while all the versions are kept in the data structure so that any previous version can still be queried. Persistent data structures aim at recording all versions accurately, which results in a space requirement that is at least linear to the number of updates. In many of today's big data applications, in particular for high-speed streaming data, the volume and velocity of the data are so high that we cannot afford to store everything. Therefore, streaming algorithms have received a lot of attention in the research community, which use only sublinear space by sacrificing slightly on accuracy. All streaming algorithms work by maintaining a small data structure in memory, which is usually called a *sketch*, *summary*, or *synopsis*. The sketch is updated upon the arrival of every element in the stream, thus is *ephemeral*, meaning that it can only answer queries about the current status of the stream. In this paper, we aim at designing persistent sketches, thereby giving streaming algorithms the ability to answer queries about the stream at any prior time.

Scalable Distributed Stream Join Processing

Qian Lin, National University of Singapore; Beng Chin Ooi, National University of Singapore; Zhengkui Wang, National University of Singapore; Cui Yu, Monmouth University

Efficient and scalable stream joins play an important role in performing real-time analytics for many cloud applications. However, like in conventional database processing, online theta-joins over data streams are computationally expensive and moreover, being memory-based processing, they impose high memory requirement on the system. In this paper, we propose a novel stream join model, called join-biclique, which organizes a large cluster as a complete bipartite graph. Join-biclique has several strengths over state-of-the-art techniques, including memory-efficiency, elasticity and scalability. These features are essential for building efficient and scalable streaming systems. Based on join-biclique, we develop a scalable distributed stream join system, BiStream, over a large-scale commodity cluster. Specifically, BiStream is designed to support efficient full-history joins, window-based joins and online data aggregation. BiStream also supports adaptive resource management

to dynamically scale out and down the system according to its application workloads. We provide both theoretical cost analysis and extensive experimental evaluations to evaluate the efficiency, elasticity and scalability of BiStream.

SCREEN: Stream Data Cleaning under Speed Constraints

Shaouxu Song, Tsinghua University; Aoqian Zhang, Tsinghua University; Jianmin Wang, Tsinghua University; Philip Yu, University of Illinois at Chicago

Stream data are often dirty, for example, owing to unreliable sensor reading, or erroneous extraction of stock prices. Most stream data cleaning approaches employ a smoothing filter, which may seriously alter the data without preserving the original information. We argue that the cleaning should avoid changing those originally correct/clean data, a.k.a. the *minimum change principle* in data cleaning. To capture the knowledge about *what is clean*, we consider the (widely existing) constraints on the speed of data changes, such as fuel consumption per hour, or daily limit of stock prices. Guided by these semantic constraints, in this paper, we propose SCREEN, the first constraint-based approach for cleaning stream data. It is notable that existing data repair techniques clean (a sequence of) data *as a whole* and fail to support stream computation. To this end, we have to relax the global optimum over the entire sequence to the local optimum in a window. Rather than the commonly observed NP-hardness of general data repairing problems, our major contributions include (1) polynomial time algorithm for global optimum, (2) linear time algorithm towards local optimum under an efficient *Median Principle*, (3) support on out-of-order arrivals of data points, and (4) adaptive window size for balancing repair accuracy and efficiency. Experiments on real datasets demonstrate that SCREEN can show significantly higher repair accuracy than the existing approaches such as smoothing.

Location-Aware Pub/Sub System: When Continuous Moving Queries Meet Dynamic Event Streams

Long Guo, National University of Singapore; Dongxiang Zhang, National University of Singapore; Guoliang Li, Tsinghua University; Kian-Lee Tan, National University of Singapore; Zhifeng Bao, RMIT University

In this paper, we propose a new location-aware pub/sub system, Elaps, that continuously monitors moving users subscribing to dynamic event streams from social media and E-commerce applications. Users are notified instantly when there is a matching event nearby. To the best of our knowledge, Elaps is the first to take into account continuous moving queries against dynamic event streams. Like existing works on continuous moving query processing, Elaps employs the concept of safe region to reduce communication overhead. However, unlike existing works which assume data from publishers are static, updates to safe regions may be triggered by newly arrived events. In Elaps, we develop a concept called *impact region* that allows us to identify whether a safe region is affected by newly arrived events. Moreover, we propose a novel cost model to optimize the safe region size to keep the communication overhead low. Based on the cost model, we design two incremental methods, iGM and idGM, for safe region construction. In addition, Elaps uses boolean expression, which is more expressive than keywords, to model user intent and we propose a novel index, BEQ-Tree, to handle spatial boolean expression matching. In our experiments, we use geo-tweets from Twitter and venues from Foursquare to simulate publishers and boolean expressions generated from AOL search log to represent users intentions. We test user movement in both synthetic trajectories and real taxi trajectories. The results show that Elaps can significantly reduce the communication overhead and disseminate events to users in real-time.

SIGMOD Research 12: Spatial data

Optimal Spatial Dominance: An Effective Search of Nearest Neighbor Candidates

Xiaoyang Wang, The University Of New South Wales; Ying Zhang, University of Technology, Sydney; Wenjie Zhang, The University Of New South Wales; Xuemin Lin, The University Of New South Wales; Muhammad Aamir Cheema, Monash University

In many domains such as computational geometry and database management, an object may be described by multiple instances (points). Then the distance (or similarity) between two objects is captured by the pair-wise distances among their instances. In the past, numerous nearest neighbor (NN) functions have been proposed to define the distance between objects with multiple instances and to identify the NN object. Nevertheless, considering that a user may not have a specific NN function in mind, it is desirable to provide her with a set of NN candidates. Ideally, the set of NN candidates must include every object that is NN for at least one of the NN functions and must exclude every non-promising object. However, no one has studied the problem of NN candidates computation from this perspective. Although some of the existing works aim at returning a set of candidate objects, they do not focus on the NN functions while computing the candidate objects. As a result, they either fail to include an NN object w.r.t. some NN functions or include a large number of unnecessary objects that have no potential to be the NN regardless of the NN functions. Motivated by this, we classify the existing NN functions for objects with multiple instances into three families by characterizing their key features. Then, we advocate three spatial dominance operators to compute NN candidates where each operator is optimal w.r.t. different coverage of NN functions. Efficient algorithms are proposed for the dominance check and corresponding NN candidates computation. Extensive empirical study on real and synthetic datasets shows that our proposed operators can significantly reduce the number of NN candidates. The comprehensive performance evaluation demonstrates the efficiency of our computation techniques.

THERMAL-JOIN: A Scalable Spatial Join for Dynamic Workloads

Farhan Tauheed, Oracle Labs; Thomas Heinis, Imperial College; Anastasia Ailamaki, Ecole Polytechnique Federale de Lausanne

Simulations have become ubiquitous in many domains of science. Today scientists study natural phenomena by first building massive three-dimensional spatial models and then by simulating the models at discrete intervals of time to mimic the behavior of natural phenomena. One frequently occurring challenge during simulations is the repeated computation of spatial self-joins of the model at each simulation time step. The join is performed to access a group of neighboring spatial objects (groups of particles, molecules or cosmological objects) so that scientists can calculate the cumulative effect (like gravitational force) on an object. Computing a self-join even in memory, soon becomes a performance bottleneck in simulation applications. The problem becomes even worse as scientists continue to improve the precision of simulations by increasing the number as well as the size (3D extent) of the objects. This leads to an exponential increase in join selectivity that challenges the performance and scalability of state-of-the-art approaches. We propose THERMAL-JOIN, a novel spatial self-join algorithm for dynamic memory-resident workloads. The algorithm groups objects in spatial proximity together into hot spots. Hot spots minimize the cost of computing join as objects assigned to a hot spot are guaranteed to overlap with each other. Using a nested spatial grid, THERMAL-JOIN partitions and indexes the dataset to locate hot spots. With experiments we show that our approach provides a speedup between 8 to 12x compared to the state of the art and also scales as scientists improve the precision of their simulations.

Indexing Metric Uncertain Data for Range Queries

Lu Chen, Zhejiang University; Yunjun Gao, Zhejiang University; Xinhan Li, Zhejiang University; Christian S. Jensen, Aalborg University; Gang Chen, Zhejiang University; Baihua Zheng, Singapore Management University

Range queries in metric spaces have applications in many areas such as multimedia retrieval, computational biology, and location-based services, where metric uncertain data exists in different forms, resulting from equipment limitations, high-throughput sequencing technologies, privacy preservation, or others. In this paper, we represent metric uncertain data by using an object-level model and a bi-level model, respectively. Two novel indexes, the uncertain pivot B+-tree (UPB-tree) and the uncertain pivot B+-forest (UPB-forest), are proposed accordingly in order to support probabilistic range queries w.r.t. a wide range of uncertain data types and similarity metrics. Both index structures use a small set of effective pivots chosen based on a newly defined criterion, and employ the B+-tree(s) as the underlying index. By design, they are easy to be integrated into any existing DBMS. In addition, we present efficient metric probabilistic range query algorithms, which utilize the validation and pruning techniques based on our derived probability lower and upper bounds. Extensive experiments with both real and synthetic data sets demonstrate that, compared against existing state-of-the-art indexes for metric uncertain data, the UPB-tree and UPB-forest incur much lower construction costs, consume smaller storage spaces, and can support more efficient metric probabilistic range queries.

Efficient Route Planning on Public Transportation Networks: A Labelling Approach

Sibo Wang, Nanyang Technological University; Wenqing Lin, A*STAR; Yi Yang, Fudan University; Xiaokui Xiao, Nanyang Technological University; Shuigeng Zhou, Fudan University

A public transportation network can often be modeled as a *timetable graph* where (i) each node represents a station; and (ii) each directed edge $\langle u, v \rangle$ is associated with a *timetable* that records the departure (resp. arrival) time of each vehicle at station u (resp. v). Several techniques have been proposed for various types of route planning on timetable graphs, e.g., retrieving the route from a node to another with the shortest travel time. These techniques, however, either provide insufficient query efficiency or incur significant space overheads. This paper presents *Timetable Labelling* (TTL), an efficient indexing technique for route planning on timetable graphs. The basic idea of TTL is to associate each node u with a set of *labels*, each of which records the shortest travel time from u to some other node v given a certain departure time from u ; such labels would then be used during query processing to improve efficiency. In addition, we propose query algorithms that enable TTL to support three popular types of route planning queries, and investigate how we reduce the space consumption of TTL with advanced preprocessing and label compression methods. By conducting an extensive set of experiments on real world datasets, we demonstrate that TTL significantly outperforms the states of the art in terms of query efficiency, while incurring moderate preprocessing and space overheads.

SIGMOD Research 13: Crowdsourcing

The Importance of Being Expert: Efficient Max-Finding in Crowdsourcing

Aris Anagnostopoulos, Sapienza University of Rome; Luca Becchetti, Sapienza University of Rome; Adriano Fazzone, Sapienza University of Rome; Ida Mele, Max Planck Institute for Informatics; Matteo Riondato, Stanford University

Crowdsourcing is a computational paradigm whose distinctive feature is the involvement of human workers in key steps of the computation. It is used successfully to address problems that would be hard or impossible to solve for machines. As we highlight in this work, the exclusive use of nonexpert individuals may prove ineffective in some cases, especially when the task at hand or the need for accurate solutions demand some degree of specialization to avoid excessive uncertainty and inconsistency in the answers. We address this limitation by proposing an approach that combines the wisdom of the crowd with the educated opinion of experts. We present a computational model for crowdsourcing that envisions two classes of workers with different expertise levels. One of its distinctive features is the adoption of the threshold error model, whose roots are in psychometrics and which we extend from previous theoretical work. Our computational model allows to evaluate the performance of crowdsourcing algorithms with respect to accuracy and cost. We use our model to develop and analyze an algorithm for approximating the best, in a broad sense, of a set of elements. The algorithm uses naïve and expert workers to find an element that is a constant-factor approximation to the best. We prove upper and lower bounds on the number of comparisons needed to solve this problem, showing that our algorithm uses expert and naïve workers optimally up to a constant factor. Finally, we evaluate our algorithm on real and synthetic datasets using the CrowdFlower crowdsourcing platform, showing that our approach is also effective in practice.

Minimizing Efforts in Validating Crowd Answers

Quoc Viet Hung Nguyen, EPFL; Chi Thang Duong, EPFL; Matthias Weidlich, Imperial College London; Karl Aberer, EPFL

In recent years, crowdsourcing has become essential in a wide range of Web applications. One of the biggest challenges of crowdsourcing is the quality of crowd answers as workers have wide-ranging levels of expertise and the worker community may contain faulty workers. Although various techniques for quality control have been proposed, a post-processing phase in which crowd answers are validated is still required. Validation is typically conducted by experts, whose availability is limited and who incur high costs. Therefore, we develop a probabilistic model that helps to identify the most beneficial validation questions in terms of both, improvement of result correctness and detection of faulty workers. Our approach allows us to guide the expert's work by collecting input on the most problematic cases, thereby achieving a set of high quality answers even if the expert does not validate the complete answer set. Our comprehensive evaluation using both real-world and synthetic datasets demonstrates that our techniques save up to 50% of expert efforts compared to baseline methods when striving for perfect result correctness. In absolute terms, for most cases, we achieve close to perfect correctness after expert input has been sought for only 20% of the questions.

iCrowd: An Adaptive Crowdsourcing Framework

Ju Fan, School of Computing, National University of Singapore; Guoliang Li, Department of Computer Science, Tsinghua University; Beng Chin Ooi, School of Computing, National University of Singapore; Kian-lee Tan, School of Computing, National University of Singapore; Jianhua Feng, Department of Computer Science, Tsinghua University

Crowdsourcing is widely accepted as a means for resolving tasks that machines are not good at. Unfortunately, Crowdsourcing may yield relatively low-quality results if there is no proper quality control. Although previous studies attempt to eliminate “bad” workers by using qualification tests, the accuracies estimated from qualifications may not be accurate, because workers have diverse accuracies across tasks. Thus, the quality of the results could be further improved by selectively assigning tasks to the workers who are well acquainted with the tasks. To this end, we propose an adaptive crowdsourcing framework, called iCrowd. iCrowd on-the-fly estimates accuracies of a worker by evaluating her performance on the completed tasks, and predicts which tasks the worker is well acquainted with. When a worker requests for a task, iCrowd assigns her a task, to which the worker has the highest estimated accuracy among all online workers. Once a worker submits an answer to a task, iCrowd analyzes her answer and adjusts estimation of her accuracies to improve subsequent task assignments. This paper studies the challenges that arise in iCrowd. The first is how to estimate diverse accuracies of a worker based on her completed tasks. The second is instant task assignment. We deploy iCrowd on Amazon Mechanical Turk, and conduct extensive experiments on real datasets. Experimental results show that iCrowd achieves higher quality than existing approaches.

QASCA: A Quality-Aware Task Assignment System for Crowdsourcing Applications

Yudian Zheng, The University of Hong Kong; Jiannan Wang, UC Berkeley; Guoliang Li, Tsinghua University; Reynold Cheng, The University of Hong Kong; Jianhua Feng, Tsinghua University

A crowdsourcing system, such as the Amazon Mechanical Turk (AMT), provides a platform for a large number of questions to be answered by Internet workers. Such systems have been shown to be useful to solve problems that are difficult for computers, including entity resolution, sentiment analysis, and image recognition. In this paper, we investigate the online task assignment problem: Given a pool of n questions, which of the k questions should be assigned to a worker? A poor assignment may not only waste time and money, but may also hurt the quality of a crowdsourcing application that depends on the workers' answers. We propose to consider quality measures (also known as evaluation metrics) that are relevant to an application during the task assignment process. Particularly, we explore how Accuracy and F-score, two widely-used evaluation metrics for crowdsourcing applications, can facilitate task assignment. Since these two metrics assume that the ground truth of a question is known, we study their variants that make use of the probability distributions derived from workers' answers. We further investigate online assignment strategies, which enables optimal task assignments. Since these algorithms are expensive, we propose solutions that attain high quality in linear time. We develop a system called the Quality-Aware Task Assignment System for Crowdsourcing Applications (QASCA) on top of AMT. We evaluate our approaches on five real crowdsourcing applications. We find that QASCA is efficient, and attains better result quality (of more than 8% improvement) compared with existing methods.

tDP: An Optimal-Latency Budget Allocation Strategy for Crowdsourced MAXIMUM Operations

Vasilis Verroios, Stanford University; Peter Lofgren, Stanford University; Hector Garcia-Molina, Stanford University

Latency is a critical factor when using a crowdsourcing platform to solve a problem like entity resolution or sorting. In practice, most frameworks attempt to reduce latency by heuristically splitting a budget of questions into rounds, so that after each round the answers are analyzed and new questions are selected. We focus on one of the most extensively studied crowdsourcing operations, the MAX operation (finding the best element in a collection under human criteria), and we study the problem of budget allocation into rounds for this operation. We provide a polynomial-time dynamic-programming budget allocation algorithm that minimizes the latency when questions form tournaments in each round. Furthermore, we study the general case where questions can be asked in any arbitrary way in each round. Our theoretical results for the general case indicate that our approach is also optimal under certain worst and average-case scenarios. We compare our approach to alternatives on Amazon Mechanical Turk, where many of our theory assumptions do not necessarily hold. We find that our approach is also optimal in practice and achieves a notable improvement over alternatives in most cases.

SIGMOD Research 14: Indexing & Performance

Cache-Efficient Aggregation: Hashing Is Sorting

Ingo Müller, Karlsruhe Institute of Technology / SAP SE; Peter Sanders, Karlsruhe Institute of Technology; Arnaud Lacurie, SAP SE; Wolfgang Lehner, Dresden University of Technology; Franz Färber, SAP SE

For decades researchers have studied the duality of hashing and sorting for the implementation of the relational operators, especially for efficient aggregation. Depending on the underlying hardware and software architecture, the specifically implemented algorithms, and the data sets used in the experiments, different authors came to different conclusions about which is the better approach. In this paper we argue that in terms of cache efficiency, the two paradigms are actually the same. We support our claim by showing that the complexity of hashing is the same as the complexity of sorting in the external memory model. Furthermore we make the similarity of the two approaches obvious by designing an algorithmic framework that allows to switch seamlessly between hashing and sorting during execution. The fact that we mix hashing and sorting routines in the same algorithmic framework allows us to leverage the advantages of both approaches and makes their similarity obvious. On a more practical note, we also show how to achieve very low constant factors by tuning both the hashing and the sorting routines to modern hardware. Since we observe a complementary dependency of the constant factors of the two routines to the locality of the input, we exploit our framework to switch to the faster routine where appropriate. The result is a novel relational aggregation algorithm that is cache-efficient---independently and without prior knowledge of input skew and output cardinality---, highly parallelizable on modern multi-core systems, and operating at a speed close to the memory bandwidth, thus outperforming the state-of-the-art by up to 3.7x.

Efficient Similarity Join and Search on Multi-Attribute Data

Guoliang Li, Tsinghua University; Jian He, Tsinghua University; Dong Deng, Tsinghua University; Jian Li, Tsinghua University

In this paper we study similarity join and search on multi-attribute data. Traditional methods on single-attribute data have pruning power only on single attributes and cannot efficiently support multi-attribute data. To address this problem, we

propose a prefix tree index which has holistic pruning ability on multiple attributes. We propose a cost model to quantify the prefix tree which can guide the prefix tree construction. Based on the prefix tree, we devise a filter-verification framework to support similarity search and join on multi-attribute data. The filter step prunes a large number of dissimilar results and identifies some candidates using the prefix tree and the verification step verifies the candidates to generate the final answer. For similarity join, we prove that constructing an optimal prefix tree is NP-complete and develop a greedy algorithm to achieve high performance. For similarity search, since one prefix tree cannot support all possible search queries, we extend the cost model to support similarity search and devise a budget-based algorithm to construct multiple high-quality prefix trees. We also devise a hybrid verification algorithm to improve the verification step. Experimental results show our method significantly outperforms baseline approaches.

Holistic Indexing in Main-memory Column-stores

Eleni Petraki, CWI; Stratos Idreos, Harvard University; Stefan Manegold, CWI

Great database systems performance relies heavily on index tuning, i.e., creating and utilizing the best indices depending on the workload. However, the complexity of the index tuning process has dramatically increased in recent years due to ad-hoc workloads and shortage of time and system resources to invest in tuning. This paper introduces holistic indexing, a new approach to automated index tuning in dynamic environments. Holistic indexing requires zero set-up and tuning effort, relying on adaptive index creation as a side-effect of query processing. Indices are created incrementally and partially; they are continuously refined as we process more and more queries. Holistic indexing takes the state-of-the-art adaptive indexing ideas a big step further by introducing the notion of a system which never stops refining the index space, taking educated decisions about which index we should incrementally refine next based on continuous knowledge acquisition about the running workload and resource utilization. When the system detects idle CPU cycles, it utilizes those extra cycles by refining the adaptive indices which are most likely to bring a benefit for future queries. Such idle CPU cycles occur when the system cannot exploit all available cores up to 100%, i.e., either because the workload is not enough to saturate the CPUs or because the current tasks performed for query processing are not easy to parallelize to the point where all available CPU power is exploited. In this paper, we present the design of holistic indexing for column-oriented database architectures and we discuss a detailed analysis against parallel versions of state-of-the-art indexing and adaptive indexing approaches. Holistic indexing is implemented in an open-source column-store DBMS. Our detailed experiments on both synthetic and standard benchmarks (TPC-H) and workloads (SkyServer) demonstrate that holistic indexing brings significant performance gains by being able to continuously refine the physical design in parallel to query processing, exploiting any idle CPU resources.

CliffGuard: A Principled Framework for Finding Robust Database Designs

Barzan Mozafari, University of Michigan, Ann Arbor; Eugene Goh, University of Michigan, Ann Arbor; Dong Yoon, University of Michigan, Ann Arbor

A fundamental problem in database systems is choosing the best physical design, i.e., a small set of auxiliary structures that enable the fastest execution of future queries. Almost all commercial databases come with designer tools that create a number of indices or materialized views (together comprising the physical design) that they exploit during query processing. Existing designers are what we call nominal; that is, they assume that their input parameters are precisely known and equal to some nominal values. For instance, since future workload is often not known a priori, it is common for these tools to optimize for past workloads in hopes that future queries and data will be similar. In practice, however, these parameters are often noisy or missing. Since nominal designers do not take the influence of such uncertainties into account, they find designs that are sub-optimal and remarkably brittle. Often, as soon as the future workload deviates from the past, their overall performance falls off a cliff, leading to customer discontent and expensive redesigns. Thus, we propose a new type of database designer that is robust against parameter uncertainties, so that overall performance degrades more gracefully when future workloads deviate from the past. Users express their risk tolerance by deciding on how much nominal optimality they are willing to trade for attaining their desired level of robustness against uncertain situations. To the best of our knowledge, this paper is the first to adopt the recent breakthroughs in the theory of robust optimization to build a practical framework for solving some of the most fundamental problems in databases, replacing today's brittle designs with a principled world of robust designs that can guarantee predictable and consistent performance.

Exploiting Correlations for Expensive Predicate Evaluation

Manas Joglekar, Stanford University; Hector Garcia-Molina, Stanford University; Aditya Parameswaran, University of Illinois at Urbana-Champaign; Christopher Ré, Stanford University

User Defined Function (UDFs) are used increasingly to augment query languages with extra, application dependent functionality. Selection queries involving UDF predicates tend to be expensive, either in terms of monetary cost or latency. In this paper, we study ways to efficiently evaluate selection queries with UDF predicates. We provide a family of techniques for processing queries at low cost while satisfying user-specified precision and recall constraints. Our techniques

are applicable to a variety of scenarios including when selection probabilities of tuples are available beforehand, when this information is available but noisy, or when no such prior information is available. We also generalize our techniques to more complex queries. Finally, we test our techniques on real datasets, and show that they achieve significant savings in UDF evaluations of up to 80%, while incurring only a small reduction in accuracy.

SIGMOD Research 15: Data Cleaning

Query-Oriented Data Cleaning with Oracles

Moria Bergman, Tel Aviv University; Tova Milo, Tel Aviv University; Slava Novgorodov, Tel Aviv University; Wang-Chiew Tan, UC Santa Cruz

As key decisions are often made based on information contained in a database, it is important for the database to be as complete and correct as possible. For this reason, many data cleaning tools have been developed to automatically resolve inconsistencies in databases. However, data cleaning tools provide only best-effort results and usually cannot eradicate all errors that may exist in a database. Even more importantly, existing data cleaning tools do not typically address the problem of determining what information is missing from a database. To overcome the limitations of existing data cleaning techniques, we present QOCO, a novel query-oriented system for cleaning data with oracles. Under this framework, incorrect (resp. missing) tuples are removed from (added to) the result of a query through edits that are applied to the underlying database, where the edits are derived by interacting with domain experts which we model as oracle crowds. We show that the problem of determining minimal interactions with oracle crowds to derive database edits for removing (adding) incorrect (missing) tuples to the result of a query is NP-hard in general and present heuristic algorithms that interact with oracle crowds. Finally, we implement our algorithms in our prototype system QOCO and show that it is effective and efficient through a comprehensive suite of experiments.

BigDancing: A System for Big Data Cleansing

Zuhair Khayyat, King Abdullah University of Science and Technology (KAUST); Ihab Ilyas, University of Waterloo; Alekh Jindal, MIT; Samuel Madden, MIT; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Jorge-Arnulfo Quiané-Ruiz, Qatar Computing Research Institute; Nan Tang, Qatar Computing Research Institute; Si Yin, Qatar Computing Research Institute

Data cleansing approaches have usually focused on detecting and fixing errors with little attention to scaling to big datasets. This presents a serious impediment since data cleansing often involves costly computations such as enumerating pairs of tuples, handling inequality joins, and dealing with user-defined functions. In this paper, we present BigDancing, a Big Data Cleansing system to tackle efficiency, scalability, and ease-of-use issues in data cleansing. The system can run on top of most common general purpose data processing platforms, ranging from DBMSs to MapReduce-like frameworks. A user-friendly programming interface allows users to express data quality rules both declaratively and procedurally, with no requirement of being aware of the underlying distributed platform. BigDancing takes these rules into a series of transformations that enable distributed computations and several optimizations, such as shared scans and specialized joins operators. Experimental results on both synthetic and real datasets show that BigDancing outperforms existing baseline systems up to more than two orders of magnitude without sacrificing the quality provided by the repair algorithms.

Data X-Ray: A Diagnostic Tool for Data Errors

Xiaolan Wang, University of Massachusetts, Amherst; Xin Dong, Google Inc.; Alexandra Meliou, University of Massachusetts, Amherst

A lot of systems and applications are data-driven, and the correctness of their operation relies heavily on the correctness of their data. While existing data cleaning techniques can be quite effective at purging datasets of errors, they disregard the fact that a lot of errors are systematic, inherent to the process that produces the data, and thus will keep occurring unless the problem is corrected at its source. In contrast to traditional data cleaning, in this paper we focus on data diagnosis: explaining where and how the errors happen in a data generative process. We develop a large-scale diagnostic framework called Data X-Ray. Our contributions are three-fold. First, we transform the diagnosis problem to the problem of finding common properties among erroneous elements, with minimal domain-specific assumptions. Second, we use Bayesian analysis to derive a cost model that implements three intuitive principles of good diagnoses. Third, we design an efficient, highly-parallelizable algorithm for performing data diagnosis on large-scale data. We evaluate our cost model and algorithm using both real-world and synthetic data, and show that our diagnostic framework produces better diagnoses and is orders of magnitude more efficient than existing techniques.

KATARA: A Data Cleaning System Powered by Knowledge Bases and Crowdsourcing

Xu Chu, University of Waterloo; John Morcos, University of Waterloo; Ihab Ilyas, University of Waterloo; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Nan Tang, Qatar Computing Research Institute; Yin Ye, Qatar Computing Research Institute

Classical approaches to clean data have relied on using integrity constraints, statistics, or machine learning. These approaches are known to be limited in the cleaning accuracy, which can usually be improved by consulting master data and involving experts to resolve ambiguity. The advent of knowledge bases (*KBs*), both general-purpose and within enterprises, and crowdsourcing marketplaces are providing yet more opportunities to achieve higher accuracy at a larger scale. We propose *KATARA*, a knowledge base and crowd powered data cleaning system that, given a table, a *KB*, and a crowd, interprets table semantics to align it with the *KB*, identifies correct and incorrect data, and generates top-*k* possible repairs for incorrect data. Experiments show that *KATARA* can be applied to various datasets and *KB*, and can efficiently annotate data and suggest possible repairs.

Crowd-Based Deduplication: An Adaptive Approach

Sibo Wang, Nanyang Technological University; Xiaokui Xiao, Nanyang Technological University; Chun-Hee Lee, Nanyang Technological University

Data deduplication stands as a building block for data integration and data cleaning. The state-of-the-art techniques focus on how to exploit crowdsourcing to improve the accuracy of deduplication. However, they either incur significant overheads on the crowd or offer inferior accuracy. This paper presents *ACD*, a new crowd-based algorithm for data deduplication. The basic idea of *ACD* is to adopt *correlation clustering* (which is a classic machine-based algorithm for data deduplication) under a crowd-based setting. We propose non-trivial techniques to reduce the time required in performing correlation clustering with the crowd, and devise methods to postprocess the results of correlation clustering for better accuracy of deduplication. With extensive experiments on the Amazon Mechanical Turk, we demonstrate that *ACD* outperforms the states of the art by offering a high precision of deduplication while incurring moderate crowdsourcing overheads.

SIGMOD Research 16: Transactions

Minimizing Commit Latency of Transactions in Geo-Replicated Data Stores

Faisal Nawab, UC Santa Barbara; Vaibhav Arora, UC Santa Barbara; Divyakant Agrawal, US Santa Barbara; Amr El Abbadi, UC Santa Barbara

Cross datacenter replication is increasingly being deployed to bring data closer to the user and to overcome datacenter outages. The extent of the influence of wide-area communication on serializable transactions is not yet clear. In this work, we derive a lower-bound on commit latency. The sum of the commit latency of any two datacenters is at least the Round-Trip Time (RTT) between them. We use the insights and lessons learned while deriving the lower-bound to develop a commit protocol, called Helios, that achieves low commit latencies. Helios actively exchanges transaction logs (history) between datacenters. The received logs are used to decide whether a transaction can commit or not. The earliest point in the received logs that is needed to commit a transaction is decided by Helios to ensure a low commit latency. As we show in the paper, Helios is theoretically able to achieve the lower-bound commit latency. Also, in a real-world deployment on five datacenters, Helios has a commit latency that is close to the optimal.

Optimizing Optimistic Concurrency Control for Tree-Structured, Log-Structured Databases

Philip Bernstein, Microsoft Research; Sudipto Das, Microsoft Research; Bailu Ding, Cornell University; Markus Pilman, ETH Zurich

Scaling-out a database system typically requires partitioning the database across multiple servers. If applications do not partition perfectly, then transactions accessing multiple partitions end up being distributed, which has well-known scalability challenges. To address them, we describe a high-performance transaction mechanism that uses optimistic concurrency control on a multi-versioned tree-structured database stored in a shared log. The system scales out by adding servers, without partitioning the database. Our solution is modeled on the Hyder architecture, published by Bernstein, Reid, and Das at CIDR 2011. We present the design and evaluation of the first full implementation of that architecture. The core of the system is a log roll-forward algorithm, called meld, that does optimistic concurrency control. Meld is inherently sequential and is therefore the main bottleneck. Our main algorithmic contributions are optimizations to meld that significantly increase transaction throughput. They use a pipelined design that parallelizes meld onto multiple threads. The slowest pipeline stage is much faster than the original meld algorithm, yielding a 3x improvement of system throughput over the original meld algorithm.

The Homeostasis Protocol: Avoiding Transaction Coordination Through Program Analysis

Sudip Roy, Google Research & Cornell University; Lucja Kot, Cornell University; Gabriel Bender, Cornell University; Bailu Ding, Cornell University; Hossein Hojjat, Cornell University; Christoph Koch, EPFL; Nate Foster, Cornell University; Johannes Gehrke, Microsoft Corp. & Cornell University

Datastores today rely on distribution and replication to achieve improved performance and fault-tolerance. But correctness of many applications depends on strong consistency properties - something that can impose substantial overheads, since it requires coordinating the behavior of multiple nodes. This paper describes a new approach to achieving strong consistency in distributed systems while minimizing communication between nodes. The key insight is to allow the state of the system to be inconsistent during execution, as long as this inconsistency is bounded and does not affect transaction correctness. In contrast to previous work, our approach uses program analysis to extract semantic information about permissible levels of inconsistency and is fully automated. We then employ a novel homeostasis protocol to allow sites to operate independently, without communicating, as long as any inconsistency is governed by appropriate treaties between the nodes. We discuss mechanisms for optimizing treaties based on workload characteristics to minimize communication, as well as a prototype implementation and experiments that demonstrate the benefits of our approach on common transactional benchmarks.

Feral Concurrency Control: An Empirical Investigation of Modern Application Integrity

Peter Bailis, UC Berkeley; Alan Fekete, University of Sydney; Michael Franklin, UC Berkeley; Ali Ghodsi, UC Berkeley; Joseph Hellerstein, UC Berkeley; Ion Stoica, UC Berkeley

The rise of data-intensive "Web 2.0" Internet services has led to a range of popular new programming frameworks that collectively embody the latest incarnation of the vision of Object-Relational Mapping (ORM) systems, albeit at unprecedented scale. In this work, we empirically investigate modern ORM-backed applications' use and disuse of database concurrency control mechanisms. Specifically, we focus our study on the common use of feral, or application-level, mechanisms for maintaining database integrity, which, across a range of ORM systems, often take the form of declarative correctness criteria, or invariants. We quantitatively analyze the use of these mechanisms in a range of open source applications written using the Ruby on Rails ORM and find that feral invariants are the most popular means of ensuring integrity (and, by usage, are over 37 times more popular than transactions). We evaluate which of these feral invariants actually ensure integrity (by usage, up to 86.9%) and which---due to concurrency errors and lack of database support---may lead to data corruption (the remainder), which we experimentally quantify. In light of these findings, we present recommendations for database system designers for better supporting these modern ORM programming patterns, thus eliminating their adverse effects on application integrity.

SIGMOD Research 17: Hardware-Aware Query Processing

Rack-Scale In-Memory Join Processing using RDMA

Claude Barthels, ETH Zurich; Simon Loesing, ETH Zurich; Gustavo Alonso, ETH Zurich; Donald Kossmann, ETH Zurich

Database systems running on a cluster of machines, i.e. rack-scale databases, are a common architecture for many large databases and data appliances. As the data movement across machines is often a significant bottleneck, these systems typically use a low-latency, high-throughput network such as InfiniBand. To achieve the necessary performance, parallel join algorithms must take advantage of the primitives provided by the network to speed up data transfer. In this paper we focus on implementing parallel in-memory joins using Remote Direct Memory Access (RDMA), a communication mechanism to transfer data directly into the memory of a remote machine. The results of this paper are, to our knowledge, the first detailed analysis of parallel hash joins using RDMA. To capture their behavior independently of the network characteristics, we develop an analytical model and test our implementation on two different types of networks. The experimental results show that the model is accurate and the resulting distributed join exhibits good performance.

Self-Tuning, GPU-Accelerated Kernel Density Models for Multidimensional Selectivity Estimation

Max Heimel, Technische Universität Berlin; Martin Kiefer, Technische Universität Berlin; Volker Markl, Technische Universität Berlin

Quickly and accurately estimating the selectivity of multidimensional predicates is a vital part of a modern relational query optimizer. The state-of-the art in this field are multidimensional histograms, which offer good estimation quality but are complex to construct and hard to maintain. Kernel Density Estimation (KDE) is an interesting alternative that does not suffer from these problems. However, existing KDE-based selectivity estimators can hardly compete with the estimation quality of state-of-the-art methods. In this paper, we substantially expand the state-of-the-art in KDE-based selectivity estimation by improving along three dimensions: First, we demonstrate how to numerically optimize a KDE model, leading to substantially improved estimates. Second, we develop methods to continuously adapt the estimator to changes in both the database and the query workload. Finally, we show how to drastically improve the performance by pushing computations

onto a GPU. We provide an implementation of our estimator and experimentally evaluate it on a variety of datasets and workloads, demonstrating that it efficiently scales up to very large model sizes, adapts itself to database changes, and typically outperforms the estimation quality of both existing Kernel Density Estimators as well as state-of-the-art multidimensional histograms.

Rethinking SIMD Vectorization for In-Memory Databases

Orestis Polychroniou, Columbia University; Arun Raghavan, Oracle Corporation; Kenneth Ross, Columbia University

Analytical databases are continuously adapting to the underlying hardware in order to saturate all sources of parallelism. At the same time, hardware evolves in multiple directions to explore different trade-offs. The MIC architecture, one such example, strays from the mainstream CPU design by packing a larger number of simpler cores per chip, relying on SIMD instructions to fill the performance gap. Databases have been attempting to utilize the SIMD capabilities of CPUs. However, mainstream CPUs have only recently adopted wider SIMD registers and more advanced instructions, since they do not rely primarily on SIMD for efficiency. In this paper, we present novel vectorized designs and implementations of database operators, based on advanced SIMD operations, such as gathers and scatters. We study selections, hash tables, and partitioning; and combine them to build sorting and joins. Our evaluation on the MIC-based Xeon Phi co-processor as well as the latest mainstream CPUs shows that our vectorization designs are up to an order of magnitude faster than the state-of-the-art scalar and vector approaches. Also, we highlight the impact of efficient vectorization on the algorithmic design of in-memory database operators, as well as the architectural design and power efficiency of hardware, by making simple cores comparably fast to complex cores. This work is applicable to CPUs and co-processors with advanced SIMD capabilities, using either many simple cores or fewer complex cores.

A Padded Encoding Scheme to Accelerate Scans by Leveraging Skew

Yinan Li, University of Wisconsin-Madison; Craig Chasseur, University of Wisconsin-Madison; Jignesh Patel, University of Wisconsin-Madison

In-memory data analytic systems that use vertical bit-parallel scan methods generally use encoding techniques. We observe that in such environments, there is an opportunity to turn skew in both the data and predicate distributions (usually a problem for query processing) into a benefit that can be leveraged to encode the column values. This paper proposes a padded encoding scheme to address this opportunity. The proposed scheme creates encodings that map common attribute values to codes that can easily be distinguished from other codes by only examining a few bits in the full code. Consequently, scans on columns stored using the padded encoding scheme can safely prune the computation without examining all the bits in the code, thereby reducing the memory bandwidth and CPU cycles that are consumed when evaluating scan queries. Our padded encoding method results in a fixed-length encoding, as fixed-length encodings are easier to manage. However, the proposed padded encoding may produce longer (fixed-length) codes than those produced by popular order-preserving encoding methods, such as dictionary-based encoding. This additional space overhead has the potential to negate the gains from early pruning of the scan computation. However, as we demonstrate empirically, the additional space overhead is generally small, and the padded encoding scheme provides significant performance improvements.

SIGMOD Research 18: Graph Propagation, Influence, Mining

GetReal: Towards Realistic Selection of Influence Maximization Strategies in Competitive Networks

Hui Li, Xidian University; Sourav S Bhowmick, Nanyang Technological University; Jiangtao Cui, Xidian University; Yunjun Gao, Zhejiang University

State-of-the-art classical influence maximization (IM) techniques are "competition-unaware" as they assume that a group (company) finds seeds (users) in a network independent of other groups who are also simultaneously interested in finding such seeds in the same network. However, in reality several groups often compete for the same market (e.g., Samsung, HTC, and Apple for the smart phone market) and hence may attempt to select seeds in the same network. This has led to increasing body of research in devising IM techniques for competitive networks. Despite the considerable progress made by these efforts toward finding seeds in a more realistic settings, unfortunately, they still make several unrealistic assumptions (e.g., a new company being aware of a rival's strategy, alternate seed selection, etc.) making their deployment impractical in real-world networks. In this paper, we propose a novel framework based on game theory to provide a more realistic solution to the IM problem in competitive networks by jettisoning these unrealistic assumptions. Specifically, we seek to find the "best" IM strategy (an algorithm or a mixture of algorithms) a group should adopt in the presence of rivals so that it can maximize its influence. As each group adopts some strategy, we model the problem as a game with each group as competitors and the expected influences under the strategies as payoffs. We propose a novel algorithm called GetReal to find each group's best solution by leveraging the competition between different groups. Specifically, it seeks to find whether

there exist a Nash Equilibrium (NE) in a game, which guarantees that there exist an "optimal" strategy for each group. Our experimental study on real-world networks demonstrates the superiority of our solution in a more realistic environment.

Influence Maximization in Near-Linear Time: A Martingale Approach

Youze Tang, Nanyang Technological University; Yanchen Shi, Nanyang Technological University; Xiaokui Xiao, Nanyang Technological University

Given a social network G and a positive integer k , the *influence maximization* problem asks for k nodes (in G) whose adoptions of a certain idea or product can trigger the largest expected number of follow-up adoptions by the remaining nodes. This problem has been extensively studied in the literature, and the state-of-the-art technique runs in $O((k+1)(n+m) \log n/\epsilon^2)$ expected time and returns a $(1-1/e-\epsilon)$ -approximate solution with at least $1 - 1/n^l$ probability. This paper presents an influence maximization algorithm that provides the same worst-case guarantees as the state of the art, but offers significantly improved empirical efficiency. The core of our algorithm is a set of estimation techniques based on *martingales*, a classic statistical tool. Those techniques not only provide accurate results with small computation overheads, but also enable our algorithm to support a larger class of information diffusion models than existing methods do. We experimentally evaluate our algorithm against the states of the art under several popular diffusion models, using real social networks with up to 1.4 billion edges. Our experimental results show that the proposed algorithm consistently outperforms the states of the art in terms of computation efficiency, and is often orders of magnitude faster.

Community Level Diffusion Extraction

Zhiting Hu, Peking University & Carnegie Mellon University; Junjie Yao, East China Normal University; Bin Cui, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University; Eric Xing, Language Technologies Institute, Carnegie Mellon University

How does online content propagate on social networks? Billions of users generate, consume, and spread tons of information every day. This unprecedented scale of dynamics becomes invaluable to reflect our zeitgeist. However, most present diffusion extraction works have only touched individual user level and cannot obtain comprehensive clues. This paper introduces a new approach, i.e., COMMUNITY LEVEL DIFFUSION (COLD), to uncover and explore temporal diffusion. We model topics and communities in a unified latent framework, and extract inter-community influence dynamics. With a well-designed multi-component model structure and a parallel inference implementation on GraphLab, the COLD method is expressive while remaining efficient. The extracted community level patterns enable diffusion exploration from a new perspective. We leverage the compact yet robust representations to develop new prediction and analysis applications. Extensive experiments on large social datasets show significant improvement in prediction accuracy. We can also find communities play very different roles in diffusion processes depending on their interest. Our method guarantees high scalability with increasing data size.

BEAR: Block Elimination Approach for Random Walk with Restart on Large Graphs

Kijung Shin, Seoul National University; Jinhong Jung, KAIST; Sael Lee, The State University of New York (SUNY) Korea; U Kang, KAIST

Given a large graph, how can we calculate the relevance between nodes fast and accurately? Random walk with restart (RWR) provides a good measure for this purpose and has been applied to diverse data mining applications including ranking, community detection, link prediction, and anomaly detection. Since calculating RWR from scratch takes long, various preprocessing methods, most of which are related to inverting adjacency matrices, have been proposed to speed up the calculation. However, these methods do not scale to large graphs because they usually produce large and dense matrices which do not fit into memory. In this paper, we propose BEAR, a fast, scalable, and accurate method for computing RWR on large graphs. BEAR comprises the preprocessing step and the query step. In the preprocessing step, BEAR reorders the adjacency matrix of a given graph so that it contains a large and easy-to-invert submatrix, and precomputes several matrices including the Schur complement of the submatrix. In the query step, BEAR computes the RWR scores for a given query node quickly using a block elimination approach with the matrices computed in the preprocessing step. Through extensive experiments, we show that BEAR significantly outperforms other state-of-the-art methods in terms of preprocessing and query speed, space efficiency, and accuracy.

The Minimum Wiener Connector Problem

Natali Ruchansky, Boston University; Francesco Bonchi, Yahoo Labs; David Garcia-Soriano, Yahoo Labs; Francesco Gullo, Yahoo Labs; Nicolas Kourtellis, Yahoo Labs

The Wiener index of a graph is the sum of all pairwise shortest-path distances between its vertices. In this paper we study the novel problem of finding a *minimum Wiener connector*: given a connected graph $G=(V, E)$ and a set $Q \subseteq V$ of query

vertices, find a subgraph of G that connects all query vertices and has minimum Wiener index. We show that our problem admits a polynomial-time (albeit impractical) exact algorithm for the special case where the number of query vertices is bounded. We show that in general the problem is NP-hard, and has no PTAS unless $P = NP$. Our main contribution is a constant-factor approximation algorithm running in time $O(|Q||E|)$. A thorough experimentation on a large variety of real-world graphs confirms that our method returns smaller and denser solutions than other methods, and does so by adding to the query set Q a small number of "important" vertices (i.e., vertices with high centrality).

SIGMOD Research 19: Social Networks

From Group Recommendations to Group Formation

Senjuti Basu Roy, University of Washington Tacoma; Laks Lakshmanan, University of British Columbia; Rui Liu, University of Washington Tacoma

There has been significant recent interest in the area of group recommendations, where, given groups of users of a recommender system, one wants to recommend top- k items to a group that maximize the satisfaction of the group members, according to a chosen semantics of group satisfaction. Examples semantics of satisfaction of a recommended itemset to a group include the so-called least misery (LM) and aggregate voting (AV). We consider the complementary problem of how to form groups such that the users in the formed groups are most satisfied with the suggested top- k recommendations. We assume that the recommendations will be generated according to one of the two group recommendation semantics -- LM or AV. Rather than assuming groups are given, or rely on ad hoc group formation dynamics, our framework allows a strategic approach for forming groups of users in order to maximize satisfaction. We show that the problem is NP-hard to solve optimally under both semantics. Furthermore, we develop two efficient algorithms for group formation under LM and show that they achieve bounded absolute error. We develop efficient heuristic algorithms for group formation under AV. We validate our results and demonstrate the scalability and effectiveness of our group formation algorithms on two large real data sets.

Real-Time Multi-Criteria Social Graph Partitioning: A Game Theoretic Approach

Nikos Armatzoglou, Hong Kong University of Science and Technology; Huy Pham, University of Southern California; Vasilis Ntranos, University of Southern California; Dimitris Papadias, Hong Kong University of Science and Technology; Cyrus Shahabi, University of Southern California

Graph partitioning has attracted considerable attention due to its high practicality for real-world applications. It is particularly relevant to social networks because it enables the grouping of users into communities for market analysis and advertising purposes. In this paper, we introduce RMGP, a type of real-time multi-criteria graph partitioning for social networks that groups the users based on their connectivity and their similarity to a set of input classes. We consider RMGP as an on-line task, which may be frequently performed for different query parameters (e.g., classes). In order to overcome the serious performance issues associated with the large social graphs found in practice, we develop solutions based on a game theoretic framework. Specifically, we consider each user as a player, whose goal is to find the class that optimizes his objective function. We propose algorithms based on best-response dynamics, analyze their properties, and show their efficiency and effectiveness on real datasets under centralized and decentralized scenarios.

Utility-Aware Social Event-Participant Planning

Jieying She, The Hong Kong University of Science and Technology; Yongxin Tong, Beihang University; Lei Chen, The Hong Kong University of Science and Technology

Online event-based social network (EBSN) platforms are becoming popular these days. An important task of managing EBSNs is to arrange proper social events to interested users. Existing approaches usually assume that each user only attends one event or ignore location information. The overall utility of such strategy is limited in real world: 1) each user may attend multiple events; 2) attending multiple events will incur spatio-temporal conflicts and travel expenses. Thus, a more intelligent EBSN platform that provides personalized event planning for each participant is desired. In this paper, we first formally define the problem of Utility-aware Social Event-participant Planning (USEP), which is proven to be NP-hard. To solve the USEP problem, we first devise a greedy-based heuristic algorithm, which performs fast under certain circumstances but has no approximation guarantee. We then present a two-step approximation framework, which not only guarantees a $1/2$ -approximation ratio but also includes a series of optimization techniques to improve its space/time efficiency. Finally, we verify the efficiency and effectiveness of the proposed methods through extensive experiments on real and synthetic datasets.

Online Video Recommendation in Sharing Community

Xiangmin Zhou, RMIT University; Lei Chen, Hong Kong University of Science and Technology; Yanchun Zhang, Victoria University; Longbing Cao, University of Technology, Sydney; Guangyan Huang, Deakin University; Chen Wang, CSIRO

The creation of sharing communities has resulted in the astonishing increasing of digital videos, and their wide applications in the domains such as entertainment, online news broadcasting etc. The improvement of these applications relies on effective solutions for social user access to video data. This fact has driven the recent research interest in social recommendation in shared communities. Although certain effort has been put into video recommendation in shared communities, the contextual information on social users has not been well exploited for effective recommendation. In this paper, we propose an approach based on the content and social information of videos for the recommendation in sharing communities. Specifically, we first exploit a robust video cuboid signature together with the Earth Mover's Distance to capture the content relevance of videos. Then, we propose to identify the social relevance of clips using the set of users belonging to a video. We fuse the content relevance and social relevance to identify the relevant videos for recommendation. Following that, we propose a novel scheme called sub-community-based approximation together with a hash-based optimization for improving the efficiency of our solution. Finally, we propose an algorithm for efficiently maintaining the social updates in dynamic shared communities. The extensive experiments are conducted to prove the high effectiveness and efficiency of our proposed video recommendation approach.

SIGMOD Research 20: Information Extraction and Record Linking

TEGRA: Table Extraction by Global Record Alignment

Xu Chu, University of Waterloo; Yeye He, Microsoft Research; Kaushik Chakrabarti, Microsoft Research; Kris Ganjam, Microsoft Research

It is well known today that pages on the Web contain a large number of content-rich relational tables. Such tables have been systematically extracted in a number of efforts to empower important applications such as table search and schema discovery. However, a significant fraction of relational tables are *not* embedded in the standard HTML table tags, and are thus difficult to extract. In particular, a large number of relational tables are known to be in a "list" form, which contains a list of clearly separated rows that are not separated into columns. In this work, we address the important problem of automatically extracting multi-column relational tables from such lists. Our key intuition lies in the simple observation that in correctly-extracted tables, values in the same column are `emph{coherent}`, both at a syntactic and at a semantic level. Using a background corpus of over 100 million tables crawled from the Web, we quantify semantic coherence based on a statistical measure of value co-occurrence in the same column from the corpus. We then model table extraction as a principled optimization problem -- we allocate tokens in each row sequentially to a fixed number of columns, such that the sum of coherence across all pairs of values in the same column is maximized. Borrowing ideas from A* search and metric distance, we develop an efficient 2-approximation algorithm. We conduct large-scale table extraction experiments using both real Web data and proprietary enterprise spreadsheet data. Our approach considerably outperforms the state-of-the-art approaches in terms of quality, achieving over 90% F-measure across many cases.

Mining Quality Phrases from Massive Text Corpora

Jialu Liu, University of Illinois at Urbana-Champaign; Jingbo Shang, University of Illinois at Urbana-Champaign; Chi Wang, Microsoft Research; Xiang Ren, University of Illinois at Urbana-Champaign; Jiawei Han, University of Illinois at Urbana-Champaign

Text data are ubiquitous and play an essential role in big data applications. However, text data are mostly unstructured. Transforming unstructured text into structured units (e.g., semantically meaningful phrases) will substantially reduce semantic ambiguity and enhance the power and efficiency at manipulating such data using database technology. Thus mining quality phrases is a critical research problem in the field of databases. In this paper, we propose a new framework that extracts quality phrases from text corpora integrated with phrasal segmentation. The framework requires only limited training but the quality of phrases so generated is close to human judgment. Moreover, the method is scalable: both computation time and required space grow linearly as corpus size increases. Our experiments on large text corpora demonstrate the quality and efficiency of the new method.

Mining Subjective Properties on the Web

Immanuel Trummer, EPFL; Alon Halevy, Google, Inc.; Hongrae Lee, Google, Inc.; Sunita Sarawagi, Google, Inc.; Rahul Gupta, Google, Inc.

Even with the recent developments in Web search of answering queries from structured data, search engines are still limited to queries with an objective answer, such as *European capitals* or *Woody Allen movies*. However, many queries are subjective, such as *safe cities*, or *cute animals*. The underlying knowledge bases of search engines do not contain answers to these queries because they do not have a ground truth. We describe the Surveyor system that mines the dominant opinion

held by authors of Web content about whether a subjective property applies to a given entity. The evidence on which Surveyor relies is statements extracted from Web text that either support the property or claim its negation. The key challenge that Surveyor faces is that simply counting the number of positive and negative statements does not suffice, because there are multiple hidden biases with which content tends to be authored on the Web. Surveyor employs a probabilistic model of how content is authored on the Web. As one example, this model accounts for correlations between the subjective property and the frequency with which it is mentioned on the Web. The parameters of the model are specialized to each property and entity type. Surveyor was able to process a large Web snapshot within a few hours, resulting in opinions for over 4-billion entity-property combinations. We selected a subset of 500 entity-property combinations and compared our results to the dominant opinion of a large number of Amazon Mechanical Turk (AMT) workers. The predictions of Surveyor match the results from AMT in 77% of all cases (and 87% for test cases where inter-worker agreement is high), significantly outperforming competing approaches.

Microblog Entity Linking with Social Temporal Context

Wen Hua, The University of Queensland; Kai Zheng, The University of Queensland; Xiaofang Zhou, The University of Queensland

Nowadays microblogging sites, such as Twitter and Chinese Sina Weibo, have established themselves as an invaluable information source, which provides a huge collection of manually-generated tweets with broad range of topics from daily life to breaking news. Entity linking is indispensable for understanding and maintaining such information, which in turn facilitates many real-world applications such as tweet clustering and classification, personalized microblog search, and so forth. However, tweets are short, informal and error-prone, rendering traditional approaches for entity linking in documents largely inapplicable. Recent work addresses this problem by utilising information from other tweets and linking entities in a batch manner. Nevertheless, the high computational complexity makes this approach infeasible for real-time applications given the high arrival rate of tweets. In this paper, we propose an efficient solution to link entities in tweets by analyzing their social and temporal context. Our proposed framework takes into consideration three features, namely entity popularity, entity recency, and user interest information embedded in social interactions to assist the entity linking task. Effective indexing structures along with incremental algorithms have also been developed to reduce the computation and maintenance costs of our approach. Experimental results based on real tweet datasets verify the effectiveness and efficiency of our proposals.

SIGMOD Research 21: RDF and SPARQL

Graph-Aware, Workload-Adaptive SPARQL Query Caching

Nikolaos Papailiou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Panagiotis Karras, Skolkovo Institute of Science and Technology (Skoltech); Nectarios Koziris, National Technical University of Athens

The pace at which data is described, queried and exchanged using the RDF specification has been ever increasing with the proliferation of Semantic Web. Minimizing SPARQL query response times has been an open issue for the plethora of RDF stores, yet SPARQL result caching techniques have not been extensively utilized. In this work we present a novel system that addresses graph-based, workload-adaptive indexing of large RDF graphs by caching SPARQL query results. At the heart of the system lies a SPARQL query canonical labelling algorithm that is used to uniquely index and reference SPARQL query graphs as well as their isomorphic forms. We integrate our canonical labelling algorithm with a dynamic programming planner in order to generate the optimal join execution plan, examining the utilization of both primitive triple indexes and cached query results. By monitoring cache requests, our system is able to identify and cache SPARQL queries that, even if not explicitly issued, greatly reduce the average response time of a workload. The proposed cache is modular in design, allowing integration with different RDF stores. Incorporating it to an open-source, distributed RDF engine that handles large scale RDF datasets, we prove that workload-adaptive caching can reduce average response times by up to two orders of magnitude and offer interactive response times for complex workloads and huge RDF datasets.

Left Bit Right: For SPARQL Join Queries with OPTIONAL Patterns (Left-outer-joins)

Medha Atre, Unaffiliated

SPARQL basic graph pattern (BGP) (a.k.a. SQL inner-join) query optimization is a well researched area. However, optimization of OPTIONAL pattern queries (a.k.a. SQL left-outer-joins) poses additional challenges, due to the restrictions on the *reordering* of left-outer-joins. The occurrence of such queries tends to be as high as 50% of the total queries (e.g., DBPedia query logs). In this paper, we present *Left Bit Right* (LBR), a technique for *well-designed* nested BGP and OPTIONAL pattern queries. Through LBR, we propose a novel method to represent such queries using a graph of *supernodes*, which is used to aggressively prune the RDF triples, with the help of compressed indexes. We also propose novel optimization strategies -- first of a kind, to the best of our knowledge -- that combine together the characteristics of

acyclicity of queries, *minimality*, and *nullification*, *best-match* operators. In this paper, we focus on OPTIONAL patterns without UNIONS or FILTERs, but we also show how UNIONS and FILTERs can be handled with our technique using a *query rewrite*. Our evaluation on RDF graphs of up to and over one billion triples, on a commodity laptop with 8 GB memory, shows that LBR can process *well-designed* low-selectivity complex queries up to 11 times faster compared to the state-of-the-art RDF column-stores as Virtuoso and MonetDB, and for highly selective queries, LBR is at par with them.

How to Build Templates for RDF Question/Answering -- An Uncertain Graph Similarity Join Approach

Weiguo Zheng, Peking University; Lei Zou, Peking University; Xiang Lian, University of Texas-Pan American; Jeffrey Xu Yu, The Chinese University of Hong Kong; Shaouxu Song, Tsinghua University; Dongyan Zhao, Peking University

A challenging task in the natural language question answering (Q/A for short) over RDF knowledge graph is how to bridge the gap between unstructured natural language questions (NLQ) and graph-structured RDF data (G). One of the effective tools is the "template", which is often used in many existing RDF Q/A systems. However, few of them study how to generate templates automatically. To the best of our knowledge, we are the first to propose a join approach for template generation. Given a workload D of SPARQL queries and a set N of natural language questions, the goal is to find some pairs $\langle q, n \rangle$, for $q \in D \wedge n \in N$, where SPARQL query q is the best match for natural language question n . These pairs provide promising hints for automatic template generation. Due to the ambiguity of the natural languages, we model the problem above as an uncertain graph join task. We propose several structural and probability pruning techniques to speed up joining. Extensive experiments over real RDF Q/A benchmark datasets confirm both the effectiveness and efficiency of our approach.

RBench: Application-Specific RDF Benchmarking

Shi Qiao, Case Western Reserve University; Z. Meral Özsoyoğlu, Case Western Reserve University

As more RDF data management systems and RDF data querying techniques emerge, RDF benchmarks providing a controllable and comparable testing environment for applications are needed. To address the needs of diverse applications, we propose an application-specific framework, called RBench, to generate RDF benchmarks. RBench takes an RDF dataset from any application as a template, and generates a set of synthetic datasets with similar characteristics including graph structure and literal labels, for the required "size scaling factor" and the "degree scaling factor". RBench analyzes several features from the given RDF dataset, and uses them to reconstruct the new benchmark graph. A flexible query load generation process is then proposed according to the design of RBench. Efficiency and usability of RBench are demonstrated via experimental results.

ALEX: Automatic Link Exploration in Linked Data

Ahmed El-Roby, University of Waterloo; Ashraf Aboulnaga, Qatar Computing Research Institute

There has recently been an increase in the number of RDF knowledge bases published on the Internet. These rich RDF data sets can be useful in answering many queries, but much more interesting queries can be answered by integrating information from different data sets. This has given rise to research on automatically linking different RDF data sets representing different knowledge bases. This is challenging due to their scale and semantic heterogeneity. Various approaches have been proposed, but there is room for improving the quality of the generated links. In this paper, we present ALEX, a system that aims at improving the quality of links between RDF data sets by using feedback provided by users on the answers to linked data queries. ALEX starts with a set of candidate links obtained using any automatic linking algorithm. ALEX utilizes user feedback to discover new links that did not exist in the set of candidate links while preserving link precision. ALEX discovers these new links by finding links that are similar to a link approved by the user through feedback on queries. ALEX uses a Monte-Carlo reinforcement learning method to learn how to explore in the space of possible links around a given link. Our experiments on real-world data sets show that ALEX is efficient and significantly improves the quality of links.

SIGMOD Research 22: Time Series & Graph Processing

k-Shape: Efficient and Accurate Clustering of Time Series

John Paparrizos, Columbia University; Luis Gravano, Columbia University

The proliferation and ubiquity of temporal data across many disciplines has generated substantial interest in the analysis and mining of time series. Clustering is one of the most popular data mining methods, not only due to its exploratory power, but also as a preprocessing step or subroutine for other techniques. In this paper, we present k-Shape, a novel algorithm for time-series clustering. k-Shape relies on a scalable iterative refinement procedure, which creates homogeneous and well-separated clusters. As its distance measure, k-Shape uses a normalized version of the cross-correlation measure in order to

consider the shapes of time series while comparing them. Based on the properties of that distance measure, we develop a method to compute cluster centroids, which are used in every iteration to update the assignment of time series to clusters. To demonstrate the robustness of k-Shape, we perform an extensive experimental evaluation of our approach against partitional, hierarchical, and spectral clustering methods, with combinations of the most competitive distance measures. k-Shape outperforms all scalable approaches in terms of accuracy. Furthermore, k-Shape also outperforms all non-scalable (and hence impractical) combinations, with one exception that achieves similar accuracy results. However, unlike k-Shape, this combination requires tuning of its distance measure and is two orders of magnitude slower than k-Shape. Overall, k-Shape emerges as a domain-independent, highly accurate, and highly efficient clustering approach for time series with broad applications.

SMiLer: A Semi-Lazy Time Series Prediction System for Sensors

Jingbo Zhou, National University of Singapore; Anthony Tung, National University of Singapore

It is useful to predict future values in time series data, for example when there are many sensors monitoring environments such as urban space. The Gaussian Process (GP) model is considered as a promising technique for this setting. However, the GP model requires too high a training cost to be tractable for large data. Though approximation methods have been proposed to improve GP's scalability, they usually can only capture global trends in the data and fail to preserve small-scale patterns, resulting in unsatisfactory performance. We propose a new method to apply the GP for sensor time series prediction. Instead of (eagerly) training GPs on entire datasets, we custom-build query-dependent GPs on small fractions of the data for each prediction request. Implementing this idea in practice at scale requires us to overcome two obstacles. On the one hand, a central challenge with such a semi-lazy learning model is the substantial model-building effort at kNN query time, which could lead to unacceptable latency. We propose a novel two-level inverted-like index to support kNN search using the DTW on the GPU, making such "just-in-time" query-dependent model construction feasible for real-time applications. On the other hand, several parameters should be tuned for each time series individually since different sensors have different data generating processes in diverse environments. Manually configuring the parameters is usually not feasible due to the large number of sensors. To address this, we devise an adaptive auto-tuning mechanism to automatically determine and dynamically adjust the parameters for each time series with little human assistance. Our method has the following strong points: (a) it can make prediction in real time without a training phase; (b) it can yield superior prediction accuracy; and (c) it can effectively estimate the analytical predictive uncertainty. To illustrate our points, we present SMiLer, a semi-lazy time series prediction system for sensors. Extensive experiments on real-world datasets demonstrate its effectiveness and efficiency. In particular, by devising a two-level inverted-like index on the GPU with an enhanced lower bound of the DTW, SMiLer accelerates the efficiency of kNN search by one order of magnitude over its baselines. The prediction accuracy of SMiLer is better than the state-of-the-art competitors (up to 10 competitors) with better estimation of predictive uncertainty.

SQLGraph: An Efficient Relational-Based Property Graph Store

Wen Sun, IBM Research - China; Achille Fokoue, IBM Watson Research Center; Kavitha Srinivas, IBM Watson Research Center; Anastasios Kementsietsidis, Google Inc.; Gang Hu, IBM Research - China; Guotong Xie, IBM Research - China

We show that existing mature, relational optimizers can be exploited with a novel schema to give better performance for property graph storage and retrieval than popular noSQL graph stores. The schema combines relational storage for adjacency information with JSON storage for vertex and edge attributes. We demonstrate that this particular schema design has benefits compared to a purely relational or purely JSON solution. The query translation mechanism translates Gremlin queries with no side effects into SQL queries so that one can leverage relational query optimizers. We also conduct an empirical evaluation of our schema design and query translation mechanism with two existing popular property graph stores. We show that our system is 2-8 times better on query performance, and 10-30 times better in throughput on 4.3 billion edge graphs compared to existing stores.

Updating Graph Indices with a One-Pass Algorithm

Dayu Yuan, Google; Prasenjit Mitra, QCRI; Huiwen Yu, Google; C. Lee Giles, Penn State University

Indices are commonly built into graph databases in order to support fast searches. Any given graph database and the distribution of queries will change over time. Therefore, the cost of processing queries using a static graph index increases because the index is built to optimize old snapshots of the database. There is growing research interest in determining how to update a graph index with the purpose of adapting to database and query changes. Updating features in a graph index is typically an NP-hard problem. In addition, because the features are chosen from a large number of frequent subgraphs, a multi-pass algorithm is not scalable to big datasets. In order to address this issue, we propose a time-efficient *one-pass algorithm* that is designed to update a graph index by scanning each frequent subgraph at most once. The algorithm replaces a feature with a new subgraph if the latter is "better" than the former one. We use the branch and bound technique to skip

subgraphs that cannot outperform any of the features in the graph index. We further use a decomposed index and reduce the space complexity from $O(|G||Q|)$ to $O(|G| + |Q|)$, where G is database graphs and Q is a query workload. Through the empirical study, we show that the one-pass algorithm is 5-100 times faster than all previous algorithms for updating graph indices. In addition, the one-pass algorithm guarantees the return of a close to optimum solution. Our experiments show that when the one-pass algorithm is used to update an index, the query-processing speed is 1-2 times faster than that of other cutting-edge indices, i.e., the FGindex and the gIndex.

SIGMOD Research 23: Advanced Query Processing

An Incremental Anytime Algorithm for Multi-Objective Query Optimization

Immanuel Trummer, EPFL; Christoph Koch, EPFL

Query plans offer diverse tradeoffs between conflicting cost metrics such as execution time, energy consumption, or execution fees in a multi-objective scenario. It is convenient for users to choose the desired cost tradeoff in an interactive process, dynamically adding constraints and finally selecting the best plan based on a continuously refined visualization of optimal cost tradeoffs. Multi-objective query optimization (MOQO) algorithms must possess specific properties to support such an interactive process: First, they must be anytime algorithms, generating multiple result plan sets of increasing quality with low latency between consecutive results. Second, they must be incremental, meaning that they avoid regenerating query plans when being invoked several times for the same query but with slightly different user constraints. We present an incremental anytime algorithm for MOQO, analyze its complexity and show that it offers an attractive tradeoff between result update frequency, single invocation time complexity, and amortized time over multiple invocations. Those properties make it suitable to be used within an interactive query optimization process. We evaluate the algorithm in comparison with prior work on TPC-H queries; our implementation is based on the Postgres database management system.

Output-sensitive Evaluation of Prioritized Skyline Queries

Niccolo Meneghetti, University at Buffalo State University of New York; Denis Mindolin, Bloomberg L.P.; Paolo Ciaccia, University of Bologna; Jan Chomicki, University at Buffalo State University of New York

Skylines assume that all attributes are equally important, as each dimension can always be traded off for another. Prioritized skylines (p-skylines) take into account non-compensatory preferences, where some dimensions are deemed more important than others, and trade-offs are constrained by the relative importance of the attributes involved. In this paper we show that querying using non-compensatory preferences is computationally efficient. We focus on preferences that are representable with p-expressions, and develop an efficient in-memory divide-and-conquer algorithm for answering p-skyline queries. Our algorithm is output-sensitive; this is very desirable in the context of preference queries, since the output is expected to be, on average, only a small fraction of the input. We prove that our method is well behaved in both the worst- and the average-case scenarios. Additionally, we develop a general framework for benchmarking p-skyline algorithms, showing how to sample prioritized preference relations uniformly, and how to highlight the effect of data correlation on performance. We conclude our study with extensive experimental results.

Learning Generalized Linear Models Over Normalized Data

Arun Kumar, University of Wisconsin-Madison; Jeffrey Naughton, University of Wisconsin-Madison; Jignesh Patel, University of Wisconsin-Madison

Enterprise data analytics is a booming area in the data management industry. Many companies are racing to develop toolkits that closely integrate statistical and machine learning techniques with data management systems. Almost all such toolkits assume that the input to a learning algorithm is a single table. However, most relational datasets are not stored as single tables due to normalization. Thus, analysts often perform key-foreign key joins before learning on the join output. This strategy of learning after joins introduces redundancy avoided by normalization, which could lead to poorer end-to-end performance and maintenance overheads due to data duplication. In this work, we take a step towards enabling and optimizing learning over joins for a common class of machine learning techniques called generalized linear models that are solved using gradient descent algorithms in an RDBMS setting. We present alternative approaches to learn over a join that are easy to implement over existing RDBMSs. We introduce a new approach named factorized learning that pushes ML computations through joins and avoids redundancy in both I/O and computations. We study the tradeoff space for all our approaches both analytically and empirically. Our results show that factorized learning is often substantially faster than the alternatives, but is not always the fastest, necessitating a cost-based approach. We also discuss extensions of all our approaches to multi-table joins as well as to Hive.

Utilizing IDs to Accelerate Incremental View Maintenance

Yannis Katsis, UC San Diego; Kian Win Ong, UC San Diego; Yannis Papakonstantinou, UC San Diego; Kevin Keliang Zhao, UC San Diego

Prior Incremental View Maintenance (IVM) algorithms specify the view tuples that need to be modified by computing diff sets, which we call tuple-based diffs since a diff set contains one diff tuple for each to-be-modified view tuple. idIVM assumes the base tables have keys and performs IVM by computing ID-based diff sets that compactly identify the to-be-modified tuples through their IDs. This work makes the following contributions: (a) An ID-based IVM system for a large subset of SQL that includes the algebraic operators selection, join, grouping and aggregation, generalized projection involving functions, antisemijoin (and therefore negation/difference) and union. The system is based on a modular approach, allowing one to extend the supported language simply by adding one algebraic operator at-a-time, along with equations describing how ID-based changes are propagated through the operator. (b) An efficient algorithm that creates an IVM plan for a given view in four passes that are polynomial in the size of the view expression. (c) A formal analysis comparing the ID-based IVM algorithm to prior IVM approaches and analytically showing when one outperforms the other. (d) An experimental comparison of the ID-based IVM algorithm to prior IVM algorithms showing the superiority of the former in common use cases.

SIGMOD Research 24: New Models

S4: Top-k Spreadsheet-Style Search for Query Discovery

Fotis Psallidas, Columbia University; Bolin Ding, Microsoft; Kaushik Chakrabarti, Microsoft; Surajit Chaudhuri, Microsoft

An enterprise information worker is often aware of a few example tuples that should be present in the output of the query. Query discovery systems have been developed to discover project-join queries that contain the given example tuples in their output. However, they require the output to exactly contain all the example tuples and do not perform any ranking. To address this limitation, we study the problem of efficiently discovering top-k project join queries which approximately contain the given example tuples in their output. We extend our algorithms to incrementally produce results as soon as the user finishes typing/modifying a cell. Our experiments on real-life and synthetic datasets show that our proposed solution is significantly more efficient compared with applying state-of-the-art algorithms.

Proactive Annotation Management in Relational Databases

Karim Ibrahim, Worcester Polytechnic Institute; Xiao Du, Worcester Polytechnic Institute; Mohamed Eltabakh, Worcester Polytechnic Institute

Annotation management and data curation has been extensively studied in the context of relational databases. However, existing annotation management techniques share a common limitation, which is that they are all passive engines, i.e., they only manage the annotations obtained from external sources such as DB admins, domain experts, and curation tools. They neither learn from the available annotations nor exploit the annotations-to-data correlations to further enhance the quality of the annotated database. Delegating such crucial and complex tasks to end-users--especially under large-scale databases and annotation sets--is clearly the wrong choice. In this paper, we propose the Nebula system, an advanced and proactive annotation management engine in relational databases. Nebula complements the state-of-art techniques in annotation management by learning from the available annotations, analyzing their content and semantics, and understanding their correlations with the data. And then, Nebula proactively discovers and recommends potentially missing annotation-to-data attachments. We propose context-aware ranking and prioritization of the discovered attachments that take into account the relationships among the data tuples and their annotations. We also propose approximation techniques and expert-enabled verification mechanisms that adaptively maintain high-accuracy predictions while minimizing the experts' involvement. Nebula is realized on top of an existing annotation management engine, and experimentally evaluated to illustrate the effectiveness of the proposed techniques, and to demonstrate the potential gain in enhancing the quality of annotated databases.

Weighted Coverage based Reviewer Assignment

Ngai Meng Kou, University of Macau; Leong Hou U, University of Macau; Nikos Mamoulis, University of Hong Kong; Zhiguo Gong, University of Macau

Peer reviewing is a standard process for assessing the quality of submissions at academic conferences and journals. A very important task in this process is the assignment of reviewers to papers. However, achieving an appropriate assignment is not easy, because all reviewers should have similar load and the subjects of the assigned papers should be consistent with the reviewers' expertise. In this paper, we propose a generalized framework for fair reviewer assignment. We first extract the domain knowledge from the reviewers' published papers and model this knowledge as a set of topics. Then, we perform a *group assignment* of reviewers to papers, which is a generalization of the classic Reviewer Assignment Problem (RAP), considering the relevance of the papers to topics as weights. We study a special case of the problem, where reviewers are to be found for just one paper (Journal Assignment Problem) and propose an exact algorithm which is fast in practice, as opposed to brute-force solutions. For the general case of having to assign multiple papers, which is too hard to be solved

exactly, we propose a greedy algorithm that achieves a $1/2$ -approximation ratio compared to the exact solution. This is a great improvement compared to the $1/3$ -approximation solution proposed in previous work for the simpler coverage-based reviewer assignment problem, where there are no weights on topics. We theoretically prove the approximation bound of our solution and experimentally show that it is superior to the current state-of-the-art.

Distributed Online Tracking

Mingwang Tang, University of Utah; Feifei Li, University of Utah; Yufei Tao, Chinese University of Hong Kong

In online tracking, an observer S receives a sequence of values, one per time instance, from a data source that is described by a function f . A tracker T wants to continuously maintain an approximation that is within an error threshold of the value $f(t)$ at any time instance t , with small communication overhead. This problem was recently formalized and studied, and a principled approach with optimal competitive ratio was proposed. This work extends the study of online tracking to a distributed setting, where a tracker T wants to track a function f that is computed from a set of functions $\{f_1, \dots, f_m\}$ from m distributed observers and respective data sources. This formulation finds numerous important and natural applications, e.g., sensor networks, distributed systems, measurement networks, and pub-sub systems. We formalize this problem and present effective online algorithms for various topologies of a distributed system/network for different aggregate functions. Experiments on large real data sets demonstrate the excellent performance of our methods in practice.

SIGMOD INDUSTRIAL PAPER ABSTRACTS

SIGMOD Industry 1: Streaming/Real-Time/Active

TencentRec: Real-time Stream Recommendation in Practice

Yanxiang Huang, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University; Bin Cui, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University; Wenyu Zhang, Tencent Inc.; Jie Jiang, Tencent Inc.; Ying Xu, Key Lab of High Confidence Software Technologies (MOE), School of EECS, Peking University

With the arrival of the big data era, opportunities as well as challenges arise in both industry and academia. As an important service in most web applications, accurate real-time recommendation in the context of big data is of high demand. Traditional recommender systems that analyze data and update models at regular time intervals cannot satisfy the requirements of modern web applications, calling for real-time recommender systems. In this paper, we tackle the "big", "real-time" and "accurate" challenges in real-time recommendation, and propose a general real-time stream recommender system built on Storm named TencentRec from three aspects, i.e., "system", "algorithm", and "data". We analyze the large amount of data streams from a wide range of applications leveraging the considerable computation ability of Storm, together with a data access component and a data storage component developed by us. To deal with various application specific demands, we have implemented several classic practical recommendation algorithms in TencentRec, including the item-based collaborative filtering, the content based, and the demographic based algorithms. Specially, we present a practical scalable item-based CF algorithm in detail, with the super characteristics such as robust to the implicit feedback problem, incremental update and real-time pruning. With the enhancement of real-time data collection and processing, we can capture the recommendation changes in real-time. We deploy the TencentRec in a series of production applications, and observe the superiority of TencentRec in providing accurate real-time recommendations for 10 billion user requests everyday.

Twitter Heron: Stream Processing at Scale

Sanjeev Kulkarni, Twitter; Nikunj Bhagat, Twitter; Masong Fu, Twitter; Vikas Kedigehalli, Twitter; Christopher Kellogg, Twitter; Sailesh Mittal, Twitter; Jignesh Patel, University of Wisconsin; Karthik Ramasamy, Twitter; Siddarth Taneja, Twitter

Storm has long served as the main platform for real-time analytics at Twitter. However, as the scale of data being processed in real-time at Twitter has increased, along with an increase in the diversity and the number of use cases, many limitations of Storm have become apparent. We need a system that scales better, has better debug-ability, has better performance, and is easier to manage – all while working in a shared cluster infrastructure. We considered various alternatives to meet these needs, and in the end concluded that we needed to build a new real-time stream data processing system. This paper presents the design and implementation of this new system, called Heron. Heron is now the de facto stream data processing engine inside Twitter, and in this paper we also share our experiences from running Heron in production. In this paper, we also provide empirical evidence demonstrating the efficiency and scalability of Heron.

Analytics in Motion - High Performance Event-Processing AND Real-Time Analytics in the Same Database

Lucas Braun, ETH Zurich; Thomas Etter, ETH Zurich; Georgios Gasparis, ETH Zurich; Martin Kaufmann, ETH Zurich; Donald Kossmann, ETH Zurich; Daniel Widmer, ETH Zurich; Aharon Avitzur, Huawei Technologies; Anthony Iliopoulos, Huawei Technologies; Eliezer Levy, Huawei Technologies; Ning Liang, Huawei Technologies

Modern data-centric flows in the telecommunications industry require real time analytical processing over a rapidly changing and large dataset. The traditional approach of separating OLTP and OLAP workloads cannot satisfy this requirement. Instead, a new class of integrated solutions for handling hybrid workloads is needed. This paper presents an industrial use case and a novel architecture that integrates key-value-based event processing and SQL-based analytical processing on the same distributed store while minimizing the total cost of ownership. Our approach combines several well-known techniques such as shared scans, delta processing, a PAX-fashioned storage layout, and an interleaving of scanning and delta merging in a completely new way. Performance experiments show that our system scales out linearly with the number of servers. For instance, our system sustains event streams of 100,000 events per second while simultaneously processing 100 ad-hoc analytical queries per second, using a cluster of 12 commodity servers. In doing so, our system meets all response time goals of our telecommunication customers; that is, 10 milliseconds per event and 100 milliseconds for an ad-hoc analytical query. Moreover, our system beats commercial competitors by a factor of 2.5 in analytical and two orders of magnitude in update performance.

Why Big Data Industrial Systems Need Rules and What We Can Do About It

Paul Suganthan G.C., University of Wisconsin-Madison; Chong Sun, Uber; Krishna Gayatri K., University of Wisconsin-Madison; Haojun Zhang, University of Wisconsin-Madison; Frank Yang, LinkedIn; Narasimhan Rampalli, @WalmartLabs; Shishir Prasad, @WalmartLabs; Esteban Arcaute, @WalmartLabs; Ganesh Krishnan, @WalmartLabs; Rohit Deep, @WalmartLabs; Vijay Raghavendra, @WalmartLabs; AnHai Doan, University of Wisconsin-Madison

Big Data industrial systems that address problems such as classification, information extraction, and entity matching very commonly use hand-crafted rules. Today, however, little is understood about the usage of such rules. In this paper we explore this issue. We discuss how these systems differ from those considered in academia. We describe default solutions, their limitations, and reasons for using rules. We show examples of extensive rule usage in industry. Contrary to popular perceptions, we show that there is a rich set of research challenges in rule generation, evaluation, execution, optimization, and maintenance. We discuss ongoing work at WalmartLabs and UW-Madison that illustrate these challenges. Our main conclusions are (1) using rules (together with techniques such as learning and crowdsourcing) is fundamental to building semantics-intensive Big Data systems, and (2) it is increasingly critical to address rule management, given the tens of thousands of rules industrial systems often manage today in an ad-hoc fashion.

SIGMOD Industry 2: Applications

Telco Churn Prediction with Big Data

Yiqing Huang, School of Computer Science and Technology, Soochow University, Suzhou 215006, China; Fangzhou Zhu, School of Computer Science and Technology, Soochow University, Suzhou 215006, China; Mingxuan Yuan, Huawei Noah's Ark Lab, Hong Kong; Ke Deng, School of Computer Science and Information Technology RMIT University, Australia; Yanhua Li, Huawei Noah's Ark Lab, Hong Kong; Bing Ni, Huawei Noah's Ark Lab, Hong Kong; Wenyuan Dai, Huawei Noah's Ark Lab, Hong Kong; Qiang Yang, Huawei Noah's Ark Lab, Hong Kong; Jia Zeng, School of Computer Science and Technology, Soochow University, Suzhou 215006, China & Huawei Noah's Ark Lab, Hong Kong

We show that telco big data can make churn prediction much easier from the 3V's perspectives: Volume, Variety, Velocity. Experimental results confirm that the prediction performance has been significantly improved by using a large volume of training data, a large variety of features from both business support systems (BSS) and operations support systems (OSS), and a high velocity of processing new coming data. We have deployed this churn prediction system in one of the biggest mobile operators in China. From millions of active customers, this system can provide a list of prepaid customers who are most likely to churn in the next month, having 0.96 precision for the top 50000 predicted churners in the list. Automatic matching retention campaigns with the targeted potential churners significantly boost their recharge rates, leading to a big business value.

The LDBC Social Network Benchmark: Interactive Workload

Orri Erling, OpenLink Software; Alex Averbuch, Neo Technology; Josep Larriba-Pey, Sparsity Technologies; Hassan Chafi, Oracle Labs; Andrey Gubichev, TU Munich; Arnau Prat, Universitat Politècnica de Catalunya; Minh-Duc Pham, VU University Amsterdam; Peter Boncz, CWI

The Linked Data Benchmark Council (LDBC) is now two years underway and has gathered strong industrial participation for its mission to establish benchmarks, and benchmarking practices for evaluating graph data management systems. The LDBC introduced a new choke-point driven methodology for developing benchmark workloads, which combines user input with input from expert systems architects, which we outline. This paper describes the LDBC Social Network Benchmark (SNB), and presents database benchmarking innovation in terms of graph query functionality tested, correlated graph generation techniques, as well as a scalable benchmark driver on a workload with complex graph dependencies. SNB has three query workloads under development: Interactive, Business Intelligence, and Graph Algorithms. We describe the SNB Interactive Workload in detail and illustrate the workload with some early results, as well as the goals for the two other workloads.

Rethinking Data-Intensive Science Using Scalable Analytics Systems

Frank Nothaft, University of California, Berkeley; Matt Massie, University of California, Berkeley; Timothy Danford, University of California, Berkeley; Zhao Zhang, University of California, Berkeley; Uri Laserson, Cloudera; Carl Yeksigian, Genomebridge; Jey Kottalam, University of California, Berkeley; Arun Ahuja, Carl Icahn School of Medicine, Mount Sinai; Jeff Hammerbacher, Carl Icahn School of Medicine, Mount Sinai; Michael Linderman, Carl Icahn School of Medicine, Mount Sinai; Michael Franklin, University of California, Berkeley; Anthony Joseph, University of California, Berkeley; David Patterson, University of California, Berkeley

"Next generation" data acquisition technologies are allowing scientists to collect exponentially more data at a lower cost. These trends are broadly impacting many scientific fields, including genomics, astronomy, and neuroscience. We can attack the problem caused by exponential data growth by applying horizontally scalable techniques from current analytics systems to accelerate scientific processing pipelines. In this paper, we describe ADAM, an example genomics pipeline that leverages the open-source Apache Spark and Parquet systems to achieve a 28x speedup over current genomics pipelines, while reducing cost by 63%. From building this system, we were able to distill a set of techniques for implementing scientific analyses efficiently using commodity "big data" systems. To demonstrate the generality of our architecture, we then implement a scalable astronomy image processing system which achieves a 2.8–8.9x improvement over the state-of-the-art MPI-based system.

QMapper for Smart Grid: Migrating SQL-based Application to Hive

Yue Wang, Institute of Computing Technology, Chinese Academy of Sciences; Yingzhong Xu, Institute of Computing Technology, Chinese Academy of Sciences; Yue Liu, Institute of Computing Technology, Chinese Academy of Sciences; Jian Chen, Zhejiang Electric Power Corporation, China; Songlin Hu, Institute of Computing Technology, Chinese Academy of Sciences

Apache Hive has been widely used by Internet companies for big data analytics applications. It can provide the capability of compiling high-level languages into efficient MapReduce workflows, which frees users from complicated and time consuming programming. The popularity of Hive and its HiveQL-compatible systems like Impala and Shark attracts attentions from traditional enterprises as well. However, enterprise big data processing systems such as Smart Grid applications often have to migrate their RDBMS-based legacy applications to Hive rather than directly writing new logic in HiveQL. Considering their differences in syntax and cost model, manual translation from SQL in RDBMS to HiveQL is very difficult, error-prone, and often leads to poor performance. In this paper, we propose QMapper, a tool for automatically translating SQL into proper HiveQL. QMapper consists of a rule-based rewriter and a cost-based optimizer. The experiments based on the TPC-H benchmark demonstrate that, compared to manually rewritten Hive queries provided by Hive contributors, QMapper dramatically reduces the query latency on average. Our real world Smart Grid application also shows its efficiency.

SIGMOD Industry 3: Novel Systems

REEF: Retainable Evaluator Execution Framework

Markus Weimer, Microsoft; Yingda Chen, Microsoft; Byung-Gon Chun, Seoul National University; Tyson Condie, UCLA; Carlo Curino, Microsoft; Chris Douglas, Microsoft; Yunseong Lee, Seoul National University; Tony Majestro, Microsoft; Dahlia Malkhi, VMware; Sergiy Matushevych, Microsoft; Brandon Myers, University of Washington; Shравan Narayanamurthy, Microsoft; Raghu Ramakrishnan, Microsoft; Sriram Rao, Microsoft; Russel Sears, PureStorage; Beysim Sezgin, Microsoft; Julia Wang, Microsoft

Resource Managers like Apache YARN have emerged as a critical layer in the cloud computing system stack, but the developer abstractions for leasing cluster resources and instantiating application logic are very low-level. This flexibility comes at a high cost in terms of developer effort, as each application must repeatedly tackle the same challenges (e.g., fault-tolerance, task scheduling and coordination) and re-implement common mechanisms (e.g., caching, bulk-data transfers). This paper presents REEF, a development framework that provides a control-plane for scheduling and coordinating task-level (data-plane) work on cluster resources obtained from a Resource Manager. REEF provides mechanisms that facilitate resource re-use for data caching, and state management abstractions that greatly ease the development of elastic data processing work-flows on cloud platforms that support a Resource Manager service. REEF is being used to develop several commercial offerings such as the Azure Stream Analytics service. Furthermore, we demonstrate REEF development of a distributed shell application, a machine learning algorithm, and a port of the CORFU-TOCS2013 system. REEF is also currently an Apache Incubator project that has attracted contributors from several institutions. <http://reef.incubator.apache.org>

Apache Tez: A Unifying Framework for Modeling and Building Data Processing Applications

Bikas Saha, Hortonworks; Hitesh Shah, Hortonworks; Siddharth Seth, Hortonworks; Gopal Vijayaraghavan, Hortonworks; Arun Murthy, Hortonworks; Carlo Curino, Microsoft

The broad success of Hadoop has led to a fast-evolving and diverse ecosystem of application engines that are building upon the YARN resource management layer. The open-source implementation of MapReduce is being slowly replaced by a collection of engines dedicated to specific verticals. This has led to growing fragmentation and repeated efforts with each new vertical engine re-implementing fundamental features (e.g. fault-tolerance, security, stragglers mitigation, etc.) from scratch. In this paper, we introduce Apache Tez, an open-source framework designed to build data-flow driven processing runtimes. Tez provides a scaffolding and library components that can be used to quickly build scalable and efficient data-flow centric engines. Central to our design is fostering component re-use, without hindering customizability of the

performance-critical data plane. This is in fact the key differentiator with respect to the previous generation of systems (e.g. Dryad, MapReduce) and even emerging ones (e.g. Spark), that provided and mandated a fixed data plane implementation. Furthermore, Tez provides native support to build runtime optimizations, such as dynamic partition pruning for Hive. Tez is deployed at Yahoo!, Microsoft Azure, LinkedIn and numerous Hortonworks customer sites, and a growing number of engines are being integrated with it. This confirms our intuition that most of the popular vertical engines can leverage a core set of building blocks. We complement qualitative accounts of real-world adoption with quantitative experimental evidence that Tez-based implementations of Hive, Pig, Spark, and Cascading on YARN outperform their original YARN implementation on popular benchmarks (TPC-DS, TPC-H) and production workloads.

Design and Implementation of the LogicBlox System

Molham Aref, LogicBlox; Balder ten Cate, LogicBlox; Todd Green, LogicBlox; Benny Kimelfeld, LogicBlox; Dan Olteanu, LogicBlox; Emir Pasalic, LogicBlox; Todd Veldhuizen, LogicBlox; Geoffrey Washburn, LogicBlox

The LogicBlox system aims to reduce the complexity of software development for modern applications which enhance and automate decision-making and enable their users to evolve their capabilities via a "self-service" model. Our perspective in this area is informed by over twenty years of experience building dozens of mission-critical enterprise applications that are in use by hundreds of large enterprises across industries such as retail, telecommunications, banking, and government. We designed and built LogicBlox to be the system we wished we had when developing those applications. In this paper, we discuss the design considerations behind the LogicBlox system and give an overview of its implementation, highlighting innovative aspects. These include: LogiQL, a unified and declarative language based on Datalog; the use of purely functional data structures; novel join processing strategies; advanced incremental maintenance and live programming facilities; a novel concurrency control scheme; and built-in support for prescriptive and predictive analytics.

Spark SQL: Relational Data Processing in Spark

Michael Armbrust, Databricks; Reynold Xin, Databricks; Cheng Lian, Databricks; Yin Huai, Databricks; Davies Liu, Databricks; Joseph Bradley, Databricks; Xiangrui Meng, Databricks; Tomer Kaftan, UC Berkeley; Michael Franklin, UC Berkeley; Ali Ghodsi, Databricks; Matei Zaharia, Databricks

Spark SQL is a new module in Apache Spark that integrates relational processing with Spark's functional programming API. Built on our experience with Shark, Spark SQL lets Spark programmers leverage the benefits of relational processing (e.g. declarative queries and optimized storage), and lets SQL users call complex analytics libraries in Spark (e.g. machine learning). Compared to previous systems, Spark SQL makes two main additions. First, it offers much tighter integration between relational and procedural processing, through a declarative DataFrame API that integrates with procedural Spark code. Second, it includes a highly extensible optimizer, Catalyst, built using features of the Scala programming language, that makes it easy to add composable rules, control code generation, and define extension points. Using Catalyst, we have built a variety of features (e.g. schema inference for JSON, machine learning types, and query federation to external databases) tailored for the complex needs of modern data analysis. We see Spark SQL as an evolution of both SQL-on-Spark and of Spark itself, offering richer APIs and optimizations while keeping the benefits of the Spark programming model.

SIGMOD Industry 4: Performance

Large-scale Predictive Analytics in Vertica: Fast Data Transfer, Distributed Model Creation, and In-database Prediction

Shreya Prasad, HP Vertica; Arash Fard, HP Vertica; Vishrut Gupta, HP Vertica; Jorge Martinez, HP Vertica; Jeff LeFevre, HP Vertica; Vincent Xu, HP Vertica; Meichun Hsu, HP Vertica; Indrajit Roy, HP Labs

A typical predictive analytics workflow will pre-process data in a database, transfer the resulting data to an external statistical tool such as R, create machine learning models in R, and then apply the model on newly arriving data. Today, this workflow is slow and cumbersome. Extracting data from databases, using ODBC connectors, can take hours on multi-gigabyte datasets. Building models on single-threaded R does not scale. Finally, it is nearly impossible to use R or other common tools, to apply models on terabytes of newly arriving data. We solve all the above challenges by integrating HP Vertica with Distributed R, a distributed framework for R. This paper presents the design of a high performance data transfer mechanism, new data-structures in Distributed R to maintain data locality with database table segments, and extensions to Vertica for saving and deploying R models. Our experiments show that data transfers from Vertica are 6x faster than using ODBC connections. Even complex predictive analysis on 100s of gigabytes of database tables can complete in minutes, and is as fast as in-memory systems like Spark running directly on a distributed file system.

Oracle Workload Intelligence

Quoc Trung Tran, Oracle USA; Konstantinos Morfonios, Oracle USA; Neoklis Polyzotis, Google, Inc.

Analyzing and understanding the characteristics of the incoming workload is crucial in unraveling trends and tuning the performance of a database system. In this work, we present Oracle Workload Intelligence (WI), a tool for workload modeling and mining, as our attempt to infer the processes that generate a given workload. WI consists of two main functionalities. First, WI derives a model that captures the main characteristics of the workload without overfitting, which makes it likely to generalize well to unseen instances of the workload. Such a model provides insights into the most frequent code paths in the application that drives the workload, and also enables optimizations inside the database system that target sequences of query statements. Second, WI can compare the models of different snapshots of the workload to detect whether the workload has changed. Such changes might indicate new trends, regressions, problems, or even security issues. We demonstrate the effectiveness of WI with an experimental study on synthetic workloads and customer-provided application benchmarks.

Purity: Building Fast, Highly-Available Enterprise Flash Storage from Commodity Components

John Colgrove, Pure Storage; John Davis, Pure Storage; John Hayes, Pure Storage; Ethan Miller, Pure Storage; Cary Sandvig, Pure Storage; Russell Sears, Pure Storage; Ari Tamches, Pure Storage; Neil Vachharajani, Pure Storage; Feng Wang, Pure Storage

Although flash storage has largely replaced hard disks in consumer class devices, enterprise workloads pose unique challenges that have slowed adoption of flash in "performance tier" storage appliances. In this paper, we describe Purity, the foundation of Pure Storage's Flash Arrays, the first all-flash enterprise storage system to support compression, deduplication, and high-availability. Purity borrows techniques from modern database and key-value storage architectures, and introduces novel storage primitives that have wide applicability to data management systems. For instance, all writes in Purity are monotonic, and deletions are handled using an atomic predicate-based tuple elision primitive. Purity's redundancy mechanisms are optimized for SSD failure modes and performance characteristics, allowing for fast recovery from component failures and lower space overhead than the best hard disk systems. We built deduplication and data compression schemes atop these primitives. Flash changes storage capacity/performance tradeoffs: unlike disk-based systems, flash deployments are rarely performance bound. A single Purity appliance can provide over 7GiB/s of throughput on 32KiB random I/Os, even through multiple device failures, and while providing asynchronous off-site replication. Typical installations have 99.9% latencies under 1ms, and production arrays average 5.4x data reduction and 99.999% availability. Purity takes advantage of storage performance increasing more rapidly than computational performance to build a simpler (with respect to engineering, installation, and management) scale-up storage appliance that supports hundreds of terabytes of highly-available, high-performance storage. The resulting performance and capacity supports many customer deployments of multiple applications, including scale-out and parallel systems, such as MongoDB and Oracle RAC, on a single Purity appliance.

On Improving User Response Times in Tableau

Pawel Terlecki, Tableau Software; Fei Xu, Tableau Software; Marianne Shaw, Tableau Software; Valeri Kim, Tableau Software; Richard Wesley, Tableau Software

The rapid increase in data volumes and complexity of applied analytical tasks poses a big challenge for visualization solutions. It is important to keep the experience highly interactive, so that users stay engaged and can perform insightful data exploration. Query processing usually dominates the cost of visualization generation. Therefore, in order to achieve acceptable response times, one needs to utilize backend capabilities to the fullest and apply techniques, such as caching or prefetching. In this paper we discuss key data processing components in Tableau: the query processor, query caches, Tableau Data Engine [1, 2] and Data Server. Furthermore, we cover recent performance improvements related to the number and quality of remote queries, broader reuse of cached data, and application of inter and intra query parallelism.

STIGMOD Industry 5: Usability

Amazon Redshift and the Case for Simpler Data Warehouses

Anurag Gupta, Amazon Web Services; Deepak Agarwal, Amazon Web Services; Derek Tan, Amazon Web Services; Jakub Kulesza, Amazon Web Services; Rahul Pathak, Amazon Web Services; Stefano Stefani, Amazon Web Services; Vidhya Srinivasan, Amazon

Web Services Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse solution that makes it simple and cost-effective to efficiently analyze large volumes of data using existing business intelligence tools. Since launching in February 2013, it has been Amazon Web Service's (AWS) fastest growing service, with many thousands of customers and many petabytes of data under management. Amazon Redshift's pace of adoption has been a surprise to many participants in the data warehousing community. While Amazon Redshift was priced disruptively at launch, available for as little as

\$1000TB/year, there are many open-source data warehousing technologies and many commercial data warehousing engines that provide free editions for development or under some usage limit. While Amazon Redshift provides a modern MPP, columnar, scale-out architecture, so too do many other data warehousing engines. And, while Amazon Redshift is available in the AWS cloud, one can build data warehouses using EC2 instances and the database engine of one's choice with either local or network-attached storage. In this paper, we discuss an oft-overlooked differentiating characteristic of Amazon Redshift – simplicity. Our goal with Amazon Redshift was not to compete with other data warehousing engines, but to compete with non-consumption. We believe the vast majority of data is collected but not analyzed. We believe, while most database vendors target larger enterprises, there is little correlation in today's economy between data set size and company size. And, we believe the models used to procure and consume analytics technology need to support experimentation and evaluation. Amazon Redshift was designed to bring data warehousing to a mass market by making it easy to buy, easy to tune and easy to manage while also being fast and cost-effective.

ShareInsights - An Unified Approach to Full-stack Data Processing

MUKUND DESHPANDE, Persistent Systems Ltd.; DHARUVA RAY, Persistent Systems Ltd.; SAMEER DIXIT, Persistent Systems Ltd.; AVADHOOT AGASTI, Persistent Systems Ltd.

The field of data analysis seeks to extract value from data for either business or scientific benefit. This field has seen a renewed interest with the advent of big data technologies and a new organizational role called data scientist. Even with the new found focus, the task of analyzing large amounts of data is still challenging and time-consuming. The essence of data analysis involves setting up data pipe-lines which consists of several operations that are chained together - starting from data collection, data quality checks, data integration, data analysis and data visualization (including the setting up of interaction paths in that visualization). In our opinion, the challenges stem from from the technology diversity at each stage of the data pipeline as well as the lack of process around the analysis. In this paper we present a platform that aims to significantly reduce the time it takes to build data pipelines. The platform attempts to achieve this in following ways. (1) Allow the user to describe the entire data pipeline with a single language and idioms - all the way from data ingestion to insight expression (via visualization and end-user interaction). (2) Provide a rich library of parts that allow users to quickly assemble a data analysis pipeline in the language. (3) Allow for a collaboration model that allows multiple users to work together on a data analysis pipeline as well as leverage and extend prior work with minimal effort. We studied the efficacy of the platform for a data hackathon competition conducted in our organization. The hackathon provided us with a way to study the impact of the approach. Rich data pipelines which traditionally took weeks to build were constructed and deployed in hours. Consequently, we believe that the complexity of designing and running the data analysis pipeline can be significantly reduced; leading to a marked improvement in the productivity of data analysts/data scientists.

SIGMOD DEMONSTRATION ABSTRACTS

SIGMOD Demonstrations, Group A

CE-Storm: Confidential Elastic Processing of Data Streams

Nick Katsipoulakis, University of Pittsburgh; Cory Thoma, University of Pittsburgh; Eric Gratta, University of Pittsburgh; Alexandros Labrinidis, University of Pittsburgh; Adam Lee, University of Pittsburgh; Panos Chrysanthis, University of Pittsburgh

Data Stream Management Systems (DSMS) are crucial for modern high-volume/high-velocity data-driven applications, necessitating a distributed approach to processing them. In addition, data providers often require certain levels of confidentiality for their data, especially in cases of user-generated data, such as those coming out of physical activity/health tracking devices (i.e., our motivating application). This demonstration will showcase Synefo, an infrastructure that enables elastic scaling of DSMS operators, and CryptStream, a framework that provides confidentiality and access controls for data streams while allowing computation on untrusted servers, fused as CE-Storm. We will demonstrate both systems working in tandem and also visualize their behavior over time under different scenarios.

IReS: Intelligent, Multi-Engine Resource Scheduler for Big Data Analytics Workflows

Katerina Doka, National Technical University of Athens; Nikolaos Papailiou, National Technical University of Athens; Dimitrios Tsoumakos, Ionian University; Christos Mantas, National Technical University of Athens; Nectarios Koziris, National Technical University of Athens

Big data analytics tools are steadily gaining ground at becoming indispensable to businesses worldwide. The complexity of the tasks they execute is ever increasing due to the surge in data and task heterogeneity. Current analytics platforms, while successful in harnessing multiple aspects of this "data deluge", bind their efficacy to a single data and compute model and often depend on proprietary systems. However, no single execution engine is suitable for all types of computation and no single data store is suitable for all types of data. To this end, we demonstrate *IReS*, the *Intelligent Resource Scheduler* for complex analytics workflows executed over multi-engine environments. Our system models the cost and performance of the

required tasks over the available platforms. *IReS* is then able to match distinct workflow parts to the execution and/or storage engine among the available ones in order to optimize with respect to a user-defined policy. During the demo, the attendees will be able to execute workflows that match real use cases and parametrize the input datasets and optimization policy. The underlying platform supports multiple compute and data engines, allowing the user to choose any subset of them. Through the inspection of the produced plan, its execution and the collection and presentation of numerous cost and performance metrics, the audience will experience first-hand how *IReS* takes advantage of heterogeneous runtimes and data stores and effectively models operator cost and performance for actual and diverse workflows.

Exploratory Keyword Search with Interactive Input

Zhifeng Bao, RMIT University; Yong Zeng, National University of Singapore; H. V. Jagadish, School of Computing; Tok Wang Ling, National University of Singapore

Due to the intrinsic ambiguity of keyword queries, users usually need to reformulate their queries multiple times to get the desired information. Even worse, users either have no way to precisely specify their search intention, or have limited domain knowledge on the data to precisely express their search intention. Moreover, they may just have a general interest to explore the data by keyword query. Therefore, our goal is to design an exploratory search paradigm that is able to bring humans more actively into the search process, in order to meet various user information needs, ranging from simple lookup to learning and understanding of the data. Besides, keyword queries against data with structure, such as XML, can run into multiple difficulties: how to identify the search target; more types of ambiguity arise as a keyword can be part of the structure as well as content of data, etc. Effectively addressing these requires solutions to multiple challenges. While some have been addressed to some extent individually, there is no previous effort to develop a comprehensive system to meet these important user needs and meet all of these challenges. Therefore, we propose a framework called ClearMap that natively supports visualized exploratory search paradigm on XML data. In particular, we offer an interactive and visualized mechanism to present the outcome of the query, enable user to explore and manipulate the underlying data to either quickly find desired information or learn the relationship among data items, as well as provide interactive suggestions when their expected results do not exist in the data. A preliminary version of ClearMap and its source code are available for try at <http://xmclearmap.comp.nus.edu.sg>.

QE3D: Interactive Visualization and Exploration of Complex, Distributed Query Plans

Daniel Scheibli, SAP SE; Christian Dinse, SAP SE; Alexander Boehm, SAP SE

QE3D is a novel query plan visualization tool that aims at providing an intuitive and holistic view of distributed query plans executed by the SAP HANA database management system. In this demonstration, we show how its interactive, three-dimensional plan representation helps to understand and quickly identify hotspots in complex, real-world scenarios.

DataXFormer: An Interactive Data Transformation Tool

John Morcos, University of Waterloo; Ziawasch Abedjan, CSAAIL MIT; Ihab Francis Ilyas, University of Waterloo; Mourad Ouzzani, Qatar Computing Research Institute; Paolo Papotti, Qatar Computing Research Institute; Michael Stonebraker, MIT CSAAIL

While syntactic transformations require the application of a formula on the input values, such as unit conversion or date format conversions, semantic transformations, such as "zip code to city", require a look-up in some reference data. We recently presented DataXFormer, a system that leverages Web tables, Web forms, and expert sourcing to cover a wide range of transformations. In this demonstration, we present the user-interaction with DataXFormer and show scenarios on how it can be used to transform data and explore the effectiveness and efficiency of several approaches for transformation discovery, leveraging about 112 million tables and online sources.

Quality-Driven Continuous Query Execution over Out-of-Order Data Streams

Yuanzhen Ji, SAP SE; Hongjin Zhou, SAP SE; Zbigniew Jerzak, SAP SE; Anisoara Nica, SAP SE; Gregor Hackenbroich, SAP SE; Christof Fetzer, Technische Universitaet Dresden

Executing continuous queries over out-of-order data streams, where tuples are not ordered according to timestamps, is challenging; because high result accuracy and low result latency are two conflicting performance metrics. Although many applications allow trading exact query results for lower latency, they still expect the produced results to meet a certain quality requirement. However, none of existing disorder handling approaches have considered minimizing the result latency while meeting user-specified requirements on the quality of query results. In this demonstration, we showcase AQ-K-slack, an adaptive, buffer-based disorder handling approach, which supports executing sliding window aggregate queries over out-of-order data streams in a quality-driven manner. By adapting techniques from the field of sampling-based approximate query processing and control theory, AQ-K-slack dynamically adjusts the input buffer size at query runtime to minimize the result latency, while respecting a user-specified threshold on relative errors in produced query results. We demonstrate a prototype stream processing system, which extends SAP Event Stream Processor with the implementation of AQ-K-slack. Through an interactive interface, the audience will learn the effect of different factors, such as the aggregate function, the

window specification, the result error threshold, and stream properties, on the latency and the accuracy of query results. Moreover, they can experience the effectiveness of AQ-K-slack in obtaining user-desired latency vs. result accuracy trade-offs, compared to naive disorder handling approaches that make extreme trade-offs. For instance, by sacrificing 1% result accuracy, our system can reduce the result latency by 80% when compared to the state of the art.

MoDisSENSE: A Distributed Spatio-Temporal and Textual Processing Platform for Social Networking Services

Ioannis Mytilinis, National Technical University of Athens; Ioannis Giannakopoulos, National Technical University of Athens; Ioannis Konstantinou, National Technical University of Athens; Katerina Doka, National Technical University of Athens; Dimitrios Tsiatsigkos, IMIS, RC Athena; Manolis Terrovitis, IMIS, RC Athena; Lampros Giampouras, Athens Technology Center S.A.; Nectarios Koziris, National Technical University of Athens

The amount of social networking data that is being produced and consumed daily is huge and it is constantly increasing. A user's digital footprint coming from social networks or mobile devices, such as comments and check-ins contains valuable information about his preferences. The collection and analysis of such footprints using also information about the users' friends and their footprints offers many opportunities in areas such as personalized search, recommendations, etc. When the size of the collected data or the complexity of the applied methods increases, traditional storage and processing systems are not enough and distributed approaches are employed. In this work, we present MoDisSENSE, an open-source distributed platform that provides personalized search for points of interest and trending events based on the user's social graph by combining spatio-textual user generated data. The system is designed with scalability in mind, it is built using a combination of latest state-of-the-art big data frameworks and its functionality is offered through easy to use mobile and web clients which support the most popular social networks. We give an overview of its architectural components and technologies and we evaluate its performance and scalability using different query types over various cluster sizes. Using the web or mobile clients, users are allowed to register themselves with their own social network credentials, perform socially enhanced queries for POIs, browse the results and explore the automatic blog creation functionality that is extracted by analyzing already collected GPS traces.

DocRicher: An Automatic Annotation System for Text Documents Using Social Media

Qiang Hu, SeSaMe Centre; Qi Liu, School of Computing; Xiaoli Wang, SeSaMe Centre; Anthony Tung, School of Computing; Shubham Goyal, SeSaMe Centre; Jisong Yang, SeSaMe Centre

We demonstrate a system, *DocRicher*, to enrich a text document with social media, that implicitly reference certain passages of it. The aim is to provide an automatic annotation interface to satisfy users' information need, without cumbersome queries to traditional search engines. The system consists of four components: text analysis, query construction, data assignment, and user feedback. Through text analysis, the system decomposes a text document into appropriate topical passages, of which each is represented using detected key phrases. By submitting combinations of these phrases as queries to social media systems, the relevant results are used to suggest new annotations, that are linked to the corresponding passages. We have built a user-friendly visualization tool for users to browse automatically recommended annotations on their reading documents. Users are either allowed to rate a recommended annotation by accepting it or not; or add a new annotation by manually highlighting texts and adding personal comments. Both these annotations are regarded as the ground truth to derive new queries for retrieving more relevant contents. We also apply data fusion to merge the query results from various contexts and retain most relevant ones.

A Demonstration of Rubato DB: A Highly Scalable NewSQL Database System for OLTP and Big Data Applications

Li-Yan Yuan, University of Alberta; Lengdong Wu, University of Alberta; Jia-Huai You, Computing Science; Yan Chi, Shanghai Shifang Software Ltd

We propose to demonstrate Rubato DB, a highly scalable NewSQL system, supporting various consistency levels from ACID to BASE for OLTP and big data applications. Rubato DB employs the staged grid architecture with a novel formula based protocol for distributed concurrency control. Our demonstration will present Rubato DB as one NewSQL database management system running on a collection of commodity servers against two of benchmark sets. The demo attendees can modify the configuration of system size, fine-tune the query workload, and visualize the performance on the fly by the graphical user interface. Attendees can experiment with various system scales, and thus grasp the potential scalability of Rubato DB, whose performance, with the increase of the number of servers used, can achieve a linear growth for both OLTP application with the strong consistency properties and key-value storage applications with the weak consistency properties.

G-OLA: Generalized On-Line Aggregation for Interactive Analysis on Big Data

Kai Zeng, University of California, Berkeley; Sameer Agarwal, Databricks Inc.; Ankur Dave, University of California, Berkeley; Michael Armbrust, Databricks Inc.; Ion Stoica, University of California, Berkeley

Nearly 15 years ago, Hellerstein, Haas and Wang proposed online aggregation (OLA), a technique that allows users to (1) observe the progress of a query by showing iteratively refined approximate answers, and (2) stop the query execution once its result achieves the desired accuracy. In this demonstration, we present G-OLA, a novel mini-batch execution model that generalizes OLA to support general OLAP queries with arbitrarily nested aggregates using efficient delta maintenance techniques. We have implemented G-OLA in FluoDB, a parallel online query execution framework that is built on top of the Spark cluster computing framework that can scale to massive data sets. We will demonstrate FluoDB on a cluster of 100 machines processing roughly 10TB of real-world session logs from a video-sharing website. Using an ad optimization and an A/B testing based scenario, we will enable users to perform real-time data analysis via web-based query consoles and dashboards.

SIGMOD Demonstrations, Group B

Thrifty: Offering Parallel Database as a Service using the Shared-Process Approach

Petrie Wong, The University of Hong Kong; Zhian He, The Hong Kong Polytechnic University; Ziqiang Feng, The Hong Kong Polytechnic University; Wenjian Xu, The Hong Kong Polytechnic University; Eric Lo, The Hong Kong Polytechnic University

Recently, Amazon has announced Redshift, a Parallel-Database-as-a-Service (PDaaS). Redshift adopts the “virtual cluster” approach to implement multi-tenancy, which has the merit of hard isolation among tenants (i.e., tenants do not interfere even when sharing resources). However, that benefit comes with poor resource utilization due to the significant redundancy incurred in the resources. In this demonstration, we present Thrifty, a Parallel-Database-as-a-Service operated using the “shared-process” approach. Compared with Redshift, each tenant in Thrifty does not occupy an exclusive amount of resource but share the database processes together, leading to better resource utilization. To avoid contention among tenants, Thrifty uses a proper cluster design, a tenant placement scheme, and a query routing mechanism to achieve soft isolation. In the demonstration, an attendee will be invited to register with Thrifty as a tenant to rent a parallel database instance. Then the attendee will be allowed to view the dashboard of a Thrifty’s administrator. Next, the attendee will be invited to control (e.g., increase) the workload of the tenant so as to see how Thrifty carries out online reconsolidation and elastic scaling.

BenchPress: Dynamic Workload Control in the OLTP-Bench Testbed

Dana Van Aken, Carnegie Mellon University; Djellel Difallah, University of Fribourg; Andrew Pavlo, Carnegie Mellon University; Carlo Curino, Microsoft Corporation; Philippe Cudre-Mauroux, University of Fribourg

Benchmarking is an essential activity when choosing database products, tuning systems, and understanding the trade-offs of the underlying engines. But the workloads available for this effort are often restrictive and non-representative of the ever changing requirements of the modern database applications. We recently introduced OLTP-Bench, an extensible testbed for benchmarking relational databases that is bundled with 15 workloads. The key features that set this framework apart is its ability to tightly control the request rate and dynamically change the transaction mixture. This allows an administrator to compose complex execution targets that recreate real system loads, and opens the doors to new research directions involving tuning for special execution patterns and multi-tenancy. In this demonstration, we highlight OLTP-Bench’s important features through the BenchPress game. It allows users to control the benchmark behavior in real time for multiple database management systems.

Demonstrating Data Near Here - Scientific Data Search

V.M. Megler, Portland State University; David Maier, Portland State University

Prior work proposed “Data Near Here” (DNH), a data search engine for scientific archives that is modeled on Internet search engines. DNH performs a periodic, asynchronous scan of each dataset in an archive, extracting lightweight features that are combined to form a dataset summary. During a search, DNH assesses the similarity of the search terms to the summary features and returns to the user, at interactive timescales, a ranked list of datasets for further exploration and analysis. We will demonstrate the search capabilities and ancillary metadata-browsing features for an archive of observational oceanographic data. While comparing search terms to complete datasets might seem ideal, interactive search speed would be impossible with archives of realistic size. We include an analysis showing that our summary-based approach gives a reasonable approximation of such a “complete dataset” similarity measure.

Slider: an Efficient Incremental Reasoner

Jules Chevalier, Université Jean Monnet; Julien Subercaze, Université Jean Monnet; Christophe Gravier, Université Jean Monnet; Frédérique Laforest, Université Jean Monnet

The Semantic Web has gained substantial momentum over the last decade. It contributes to the manifestation of knowledge from data, and leverages implicit knowledge through reasoning algorithms. The main drawbacks of current reasoning methods over ontologies are two-fold: first they struggle to provide scalability for large datasets, and second, the batch processing reasoners who provide the best scalability so far are unable to infer knowledge from evolving data. We

contribute to solving these problems by introducing Slider, an efficient incremental reasoner. Slider goes a significant step beyond existing system, including i) performance, by more than a 70% improvement in average compared to the fastest reasoner available to the best of our knowledge, and ii) inferences on streams of semantic data, by using intrinsic features that are themselves streams-oriented. Slider is fragment agnostic and conceived to handle expanding data with a growing background knowledge base. It natively supports pdf and RDFS, and its architecture allows to extend it to more complex fragments with a minimal effort. In this demo a web-based interface allows the users to visualize the internal behaviour of Slider during the inference, to better understand its design and principles.

WANalytics: Geo-Distributed Analytics for a Data Intensive World

Ashish Vulimiri, UIUC; Carlo Curino, Microsoft; Philip Godfrey, UIUC; Thomas Jungblut, Microsoft; Konstantinos Karanasos, Microsoft; Jitendra Padhye, Microsoft; George Varghese, Microsoft

Many large organizations collect massive volumes of data each day in a geographically distributed fashion, at data centers around the globe. Despite their geographically diverse origin the data must be processed and analyzed as a whole to extract insight. We call the problem of supporting large-scale geo-distributed analytics *Wide-Area Big Data (WABD)*. To the best of our knowledge, WABD is currently addressed by copying all the data to a central data center where the analytics are run. This approach consumes expensive cross-data center bandwidth and is incompatible with data sovereignty restrictions that are starting to take shape. We instead propose *WANalytics*, a system that solves the WABD problem by orchestrating distributed query execution and adjusting data replication across data centers in order to minimize bandwidth usage, while respecting sovereignty requirements. *WANalytics* achieves an up to 360X reduction in data transfer cost when compared to the centralized approach on both real Microsoft production workloads and standard synthetic benchmarks, including TPC-CH and Berkeley Big-Data. In this demonstration, attendees will interact with a live geo-scale multi-data center deployment of *WANalytics*, allowing them to experience the data transfer reduction our system achieves, and to explore how it dynamically adapts execution strategy in response to changes in the workload and environment.

FTT: a System for Finding and Tracking Tourists in Public Transport Services

Huayu Wu, Institute for Infocomm Research; Jo-Anne Tan, Institute for Infocomm Research; Wee Siong Ng, Institute for Infocomm Research; Mingqiang Xue, Institute for Infocomm Research; Wei Chen, Institute for Infocomm Research

The tourism industry is a key economic driver for many cities. To understand tourists' traveling patterns can help both public and private relevant sectors design and improve their services to serve tourists better and get additional values from it. The existing approaches to discover tourists' traveling pattern focus on small sets of known tourists extracted from social media or other channels. The accuracy of the mining result cannot be guaranteed due to the small and bias set of samples. In this paper, we present our system FTT (Finding and Tracking Tourists) to identify tourists from public transport commuters in a city, and to further track their movements from one place to another. Our target is a large set of tourists and their trajectories extracted from public transport riding records, which more accurately represent the movements of general tourists. In particular, we design an iterative learning algorithm to find the tourists among public transport commuters, and provide interface to answer user queries on tourists' traveling patterns. The result will be visualized on top of a city map.

SharkDB: An In-Memory Storage System for Massive Trajectory Data

Haozhou Wang, The University of Queensland; Kai Zheng, The University of Queensland; Xiaofang Zhou, The University of Queensland; Shazia Sadiq, The University of Queensland

An increasing amount of motion history data, which is called trajectory, is being collected from different sources such as GPS-enabled mobile devices, surveillance cameras and social networks. However it is hard to store and manage trajectory data in traditional database systems, since its variable lengths and asynchronous sampling rates do not fit disk-based and tuple-oriented structures, which are the fundamental structures of traditional database systems. We implement a novel trajectory storage system that is motivated by the success of column store and recent development of in-memory based databases. In this storage design, we try to explore the potential opportunities, which can boost the performance of query processing for trajectory data. To achieve this, we partition the trajectories into frames as column-oriented storage in order to store the sample points of a moving object, which are aligned by the time interval, within the main memory. Furthermore, the frames can be highly compressed and well structured to increase the memory utilization ratio and reduce the CPU-cache missing. It is also easier for parallelizing data processing on the multi-core server since the frames are mutually independent.

Ringo: Interactive Graph Analytics on Big-Memory Machines

Yonathan Perez, Stanford University; Rok Sosic, Stanford University; Arijit Banerjee, Stanford University; Rohan Puttagunta, Stanford University; Martin Raison, Stanford University; Pararth Shah, Stanford University; Jure Leskovec, Stanford University

We present Ringo, a system for analysis of large graphs. Graphs provide a way to represent and analyze systems of interacting objects (people, proteins, webpages) with edges between the objects denoting interactions (friendships, physical

interactions, links). Mining graphs provides valuable insights about individual objects as well as the relationships among them. In building Ringo, we take advantage of the fact that machines with large memory and many cores are widely available and also relatively affordable. This allows us to build an easy-to-use interactive high-performance graph analytics system. Graphs also need to be built from input data, which often resides in the form of relational tables. Thus, Ringo provides rich functionality for manipulating raw input data tables into various kinds of graphs. Furthermore, Ringo also provides over 200 graph analytics functions that can then be applied to constructed graphs. We show that a single big-memory machine provides a very attractive platform for performing analytics on all but the largest graphs as it offers excellent performance and ease of use as compared to alternative approaches. With Ringo, we also demonstrate how to integrate graph analytics with an iterative process of trial-and-error data exploration and rapid experimentation, common in data mining workloads.

STORM: Spatio-Temporal Online Reasoning and Management of Large Spatio-Temporal Data

Robert Christensen, University of Utah; Feifei Li, University of Utah; Lu Wang, Hong Kong University of Science and Technology; Ke Yi, Hong Kong; Jun Tang, University of Utah; Natalee Villa, University of Utah

We present the STORM system to enable spatio-temporal online reasoning and management of large spatio-temporal data. STORM supports interactive spatio-temporal analytics through novel spatial online sampling techniques. Online spatio-temporal aggregation and analytics are then derived based on the online samples, where approximate answers with approximation quality guarantees can be provided immediately from the start of query execution. The quality of these online approximations improve over time. This demonstration proposal describes key ideas in the design of the STORM system, and presents the demonstration plan.

PAXQuery: Parallel Analytical XML Processing

Jesús Camacho-Rodríguez, Hortonworks Inc.; Dario Colazzo, Université Paris-Dauphine; Ioana Manolescu, INRIA & Université Paris-Sud; Juan Naranjo, INRIA & Université Paris-Sud

XQuery is a general-purpose programming language for processing semi-structured data, and as such, it is very expressive. As a consequence, optimizing and parallelizing complex analytics XQuery queries is still an open, challenging problem. We demonstrate PAXQuery, a novel system that parallelizes the execution of XQuery queries over large collections of XML documents. PAXQuery compiles a rich subset of XQuery into plans expressed in the PArallelization ConTracts (PACT) programming model. Thanks to this translation, the resulting plans are optimized and executed in a massively parallel fashion by the Apache Flink system. The result is a scalable system capable of querying massive amounts of XML data very efficiently, as proved by the experimental results we outline.

SIGMOD Demonstrations, Group C

Graft: A Debugging Tool For Apache Giraph

Semih Salihoglu, Stanford University; Jaeho Shin, Stanford University; Vikesh Khanna, Stanford University; Ba Quan Truong, Stanford University; Jennifer Widom, Stanford University

We address the problem of debugging programs written for Pregel-like systems. After interviewing Giraph and GPS users, we developed Graft. Graft supports the debugging cycle that users typically go through: (1) Users describe programmatically the set of vertices they are interested in inspecting. During execution, Graft captures the context information of these vertices across supersteps. (2) Using Graft's GUI, users visualize how the values and messages of the captured vertices change from superstep to superstep, narrowing in suspicious vertices and supersteps. (3) Users replay the exact lines of the `vertex.compute()` function that executed for the suspicious vertices and supersteps, by copying code that Graft generates into their development environments' line-by-line debuggers. Graft also has features to construct end-to-end tests for Giraph programs. Graft is open-source and fully integrated into Apache Giraph's main code base.

Even Metadata is Getting Big: Annotation Summarization using InsightNotes

Dongqing Xiao, Worcester Polytechnic Institute; Armir Bashllari, Worcester Polytechnic Institute; Tyler Menard, Worcester Polytechnic Institute; Mohamed Eltabakh, Worcester Polytechnic Institute

In this paper, we demonstrate the InsightNotes system, a summary-based annotation management engine over relational databases. InsightNotes addresses the unique challenges that arise in modern applications (especially scientific applications) that rely on rich and large-scale repositories of curation and annotation information. In these applications, the number and size of the raw annotations may grow beyond what end-users and scientists can comprehend and analyze. InsightNotes overcomes these limitations by integrating mining and summarization techniques with the annotation management engine in novel ways. The objective is to create concise and meaningful representations of the raw annotations, called "annotation summaries", to be the basic unit of processing. The core functionalities of InsightNotes include: (1) Extensibility, where domain experts can define the summary types suitable for their application, (2) Incremental Maintenance, where the system efficiently maintains the annotation summaries under the continuous addition of new annotations, (3) Summary-Aware

Query Processing and Propagation, where the execution engine and query operators are extended for manipulating and propagating the annotation summaries within the query pipeline under complex transformations, and (4) Zoom-in Query Processing, where end-users can interactively expand specific annotation summaries of interest and retrieve their detailed (raw) annotations. We will demonstrate the InsightNotes's features using a real-world annotated database from the ornithological domain (the science of studying birds). We will design an interactive demonstration that engage the audience in annotating the data, visualizing how annotations are summarized and propagated, and zooming-in when desired to retrieve more details.

StoryPivot: Comparing and Contrasting Story Evolution

Anja Gruenheid, ETH Zurich; Donald Kossmann, ETH Zurich; Theodoros Rekatsinas, University of Maryland; Divesh Srivastava, AT&T Labs-Research

As the world evolves around us, so does the digital coverage of it. Events of diverse types, associated with different actors and various locations, are continuously captured by multiple information sources such as news articles, blogs, social media etc. day by day. In the digital world, these events are represented through information snippets that contain information on the involved entities, a description of the event, when the event occurred, etc. In our work, we observe that events (and their corresponding digital representations) are often inter-connected, i.e., they form stories which represent evolving relationships between events over time. Take as an example the plane crash in Ukraine in July 2014 which involved multiple entities such as "Ukraine", "Malaysia", and "Russia" and multiple events ranging from the actual crash to the incident investigation and the presentation of the investigator's findings. In this demonstration we present StoryPivot, a framework that helps its users to detect evolving stories in event datasets over time. To resolve stories, we differentiate between story identification, the problem of connecting events over time within a source, and story alignment, the problem of integrating stories across sources. The goal of this demonstration is to present an interactive exploration of both these problems and how events can be dynamically interpreted and put into context in real-world datasets.

The Flatter, the Better – Query Compilation Based on the Flattening Transformation

Alexander Ulrich, Universität Tübingen; Torsten Grust, Universität Tübingen

We demonstrate the insides and outs of a query compiler based on the flattening transformation, a translation technique designed by the programming language community to derive efficient data-parallel implementations from iterative programs. Flattening admits the straightforward formulation of intricate query logic—including deeply nested loops over (possibly ordered) data or the construction of rich data structures. To demonstrate the level of expressiveness that can be achieved, we will bring a compiler frontend that accepts queries embedded into the Haskell programming language. Compilation via flattening takes places in a series of simple steps all of which will be made tangible by the demonstration. The final output is a program of lifted primitive operations which existing query engines can efficiently implement. We provide backends based on PostgreSQL and VectorWise to make this point—however, most set-oriented or data-parallel engines could benefit from a flattening-based query compiler.

D2WORM: A Management Infrastructure for Distributed Data-centric Workflows

Martin Jergler, Technische Universität München; Mohammad Sadoghi, IBM; Hans-Arno Jacobsen, Technische Universität München

Unlike traditional activity-flow-based models, data-centric workflows primarily focus on the data to drive a business. This enables the unification of operational management, concurrent process analytics, compliance with process or associated data constraints, and adaptability to changing environments. In this demonstration, we present D2Worm, a Distributed Data-centric Workflow Management system. D2Worm allows users to (1) graphically model data-centric workflows in a declarative fashion based on the Guard-Stage-Milestone (GSM) meta-model, (2) automatically compile the modelled workflow into several fine-granular workflow units (WFUs), and (3) deploy these WFUs on distributed infrastructures. A WFU is a system component that manages a subset of the workflow's data model and, at the same time, represents part of the global control flow by evaluating conditions over the data. WFUs communicate with each other over a publish/subscribe messaging infrastructure that allows the architecture to scale from a single node to dozens of machines distributed over different data-centers. In addition, D2Worm is able to (4) concurrently execute multiple workflow instances and monitor their behavior in real-time.

NL2CM: A Natural Language Interface to Crowd Mining

Yael Amerdamer, Tel Aviv University; Anna Kukliansky, Tel Aviv University; Tova Milo, Tel Aviv University

The joint processing of general data, which can refer to objective data such as geographical locations, with individual data, which is related to the habits and opinions of individuals, is required in many real-life scenarios. For this purpose, crowd mining platforms combine searching knowledge bases for general data, with mining the crowd for individual, unrecorded data. Existing such platforms require queries to be stated in a formal language. To bridge the gap between naive users, who are not familiar with formal query languages, and crowd mining platforms, we develop NL2CM, a prototype system which

translates natural language (NL) questions into well-formed crowd mining queries. The mix of general and individual information needs raises unique challenges. In particular, the different types of needs must be identified and translated into separate query parts. To account for these challenges, we develop new, dedicated modules and embed them within the modular and easily extensible architecture of NL2CM. Some of the modules interact with the user during the translation process to resolve uncertainties and complete missing data. We demonstrate NL2CM by translating questions of the audience, in different domains, into NL2CM, a crowd mining query language which is based on SPARQL.

Optimistic Recovery for Iterative Dataflows in Action

Sergey Dudoladov, Technische Universität Berlin; Chen Xu, Technische Universität Berlin; Sebastian Schelter, Technische Universität Berlin; Asterios Katsifodimos, Technische Universität Berlin; Stephan Ewen, Data Artisans GmbH; Kostas Tzoumas, Data Artisans GmbH; Volker Markl, Technische Universität Berlin

Over the past years, parallel dataflow systems have been employed for advanced analytics in the field of data mining where many algorithms are iterative. These systems typically provide fault tolerance by periodically checkpointing the algorithm's state and, in case of failure, restoring a consistent state from a checkpoint. In prior work, we presented an optimistic recovery mechanism that in certain cases eliminates the need to checkpoint the intermediate state of an iterative algorithm. In case of failure, our mechanism uses a compensation function to transit the algorithm to a consistent state, from which the execution can continue and successfully converge. Since this recovery mechanism does not checkpoint any state, it achieves optimal failure-free performance while guaranteeing fault tolerance. In this paper, we demonstrate our recovery mechanism with the Apache Flink data processing engine. During our demonstration, attendees will be able to run graph algorithms and trigger failures to observe the algorithms recovering with compensation functions instead of checkpoints.

A Secure Search Engine for the Personal Cloud

Saliha Lallali, INRIA Rocquencourt; Nicolas Ancaix, INRIA Rocquencourt; Iulian Sandu Popa, INRIA Rocquencourt; Philippe Pucheral, INRIA Rocquencourt

The emerging Personal Cloud paradigm holds the promise of a Privacy-by-Design storage and computing platform where personal data remain under the individual's control while being shared by valuable applications. However, leaving the data management control to user's hands pushes the security issues to the user's platform. This demonstration presents a Secure Personal Cloud Platform relying on a query and access control engine embedded in a tamper resistant hardware device connected to the user's platform. The main difficulty lies in the design of an inverted document index and its related search and update algorithms capable of tackling the strong hardware constraints of these devices. We have implemented our engine on a real tamper resistant hardware device and present its capacity to regulate the access to a personal dataspace. The objective of this demonstration is to show (1) that secure hardware is a key enabler of the Personal Cloud paradigm and (2) that new embedded indexing and querying techniques can tackle the hardware constraints of tamper-resistant devices and provide scalable solutions for the Personal Cloud.

A SQL Debugger Built from Spare Parts -- Turning a SQL:1999 Database System into Its Own Debugger

Benjamin Dietrich, Universität Tübingen; Torsten Grust, Universität Tübingen

We demonstrate a new incarnation of Habitat, an observational debugger for SQL. In observational debugging, users highlight parts of a—presumably faulty—query to observe the evaluation of SQL subexpressions and learn about the query's actual runtime behavior. The present version of Habitat has been redesigned from scratch and employs a query instrumentation technique that exclusively relies on the SQL facilities of the underlying RDBMS. We particularly shed light on new features like (1) the debugging of recursive SQL queries and (2) the observation of row groups (before and after aggregation). Habitat can turn any reasonably modern SQL:1999 RDBMS into its own language-level SQL debugger.

Just can't get enough - Synthesizing Big Data

Tilmann Rabl, University of Toronto; Manuel Danisch, bankmark; Michael Frank, bankmark; Sebastian Schindler, bankmark; Hans-Arno Jacobsen, University of Toronto

With the rapidly decreasing prices for storage and storage systems ever larger data sets become economical. While only few years ago only successful transactions would be recorded in sales systems, today every user interaction will be stored for ever deeper analysis and richer user modeling. This has led to the development of big data systems, which offer high scalability and novel forms of analysis. Due to the rapid development and ever increasing variety of the big data landscape, there is a pressing need for tools for testing and benchmarking. Vendors have little options to showcase the performance of their systems but to use trivial data sets like TeraSort or WordCount. Since customers' real data is typically subject to privacy regulations and rarely can be utilized, simplistic proof-of-concepts have to be used, leaving both, customers and vendors, unclear of the target use-case performance. As a solution, we present an automatic approach to data synthetization from existing data sources. Our system enables a fully automatic generation of large amounts of complex, realistic, synthetic data.

PODS PAPER ABSTRACTS

PODS Session 1: Repairs and Ontologies

Dichotomies in the Complexity of Preferred Repairs

Ronald Fagin, IBM Research; Benny Kimelfeld, LogicBlox; Phokion Kolaitis, University of California Santa Cruz.

The framework of database repairs provides a principled approach to managing inconsistencies in databases. Informally, a repair of an inconsistency database is a consistent database that differs from the inconsistent one in a “minimal way”. A fundamental problem in this framework is the repair-checking problem: given two instances, is the second a repair of the first? Here, all repairs are taken into account, and they are treated on a par with each other. There are situations, however, in which it is natural and desired to prefer one repair over another; for example, one data source is regarded to be more reliable than another, or timestamp information implies that a more recent fact should be preferred over an earlier one. Motivated by these considerations, Staworko, Chomicki and Marcinkowski introduced the framework of preferred repairs. The main characteristic of this framework is that it uses a priority relation between conflicting facts of an inconsistent database to define notions of preferred repairs. In this paper we focus on the globally-optimal repairs, in the case where the constraints are functional dependencies. Intuitively, a globally-optimal repair is a repair that cannot be improved by exchanging facts with preferred facts. In this setting, it is known that there is a fixed schema (i.e., signature and functional dependencies) where globally-optimal repair-checking is coNP-complete. Our main result is a dichotomy in complexity: for each fixed relational signature and each fixed set of functional dependencies, the globally-optimal repair-checking problem either is solvable in polynomial time or is coNP-complete. Specifically, the problem is solvable in polynomial time if for each relation symbol in the signature, the functional dependencies are equivalent to either a single functional dependency or to a set of two key constraints; in all other cases, the globally-optimal repair-checking problem is coNP-complete. We also show that there is a polynomial-time algorithm for distinguishing between the tractable and the intractable cases. The setup of preferred repairs assumes that preferences are only between conflicting facts. In the last part of the paper, we investigate the effect of this assumption on the complexity of globally-optimal repair checking. With this assumption relaxed, we give another dichotomy theorem and another polynomial-time distinguishing algorithm. Interestingly, the two dichotomies turn out to have quite different conditions for distinguishing tractability from intractability.

The Data Complexity of Consistent Query Answering for Self-Join-Free Conjunctive Queries Under Primary Key Constraints

Paraschos Koutris, University of Washington; Jef Wijsen, University of Mons.

A relational database is said to be uncertain if primary key constraints can possibly be violated. A repair (or possible world) of an uncertain database is obtained by selecting a maximal number of tuples without ever selecting two distinct tuples with the same primary key value. For any Boolean query q , $\text{CERTAINTY}(q)$ is the problem that takes an uncertain database db as input, and asks whether q is true in every repair of db . The complexity of this problem has been particularly studied for q ranging over the class of self-join-free Boolean conjunctive queries. A research challenge is to determine, given q , whether $\text{CERTAINTY}(q)$ belongs to complexity classes FO, P, or coNP-complete. In this paper, we combine existing techniques for studying the above complexity classification task. We show that for any self-join-free Boolean conjunctive query q , it can be decided whether or not $\text{CERTAINTY}(q)$ is in FO. Further, for any self-join-free Boolean conjunctive query q , $\text{CERTAINTY}(q)$ is either in P or coNP-complete, and the complexity dichotomy is effective. This settles a research question that has been open for ten years.

High-Level Why-Not Explanations using Ontologies

Balder Ten Cate, LogicBlox; Cristina Civili, Sapienza University of Rome; Evgeny Sherkhonov, University of Amsterdam; Wang-Chiew Tan, University of California Santa Cruz.

We propose a novel foundational framework for why-not explanations, that is, explanations for why a tuple is missing from a query result. Our why-not explanations leverage concepts from an ontology to provide high-level and meaningful reasons for why a tuple is missing from the result of a query. A key algorithmic problem in our framework is that of computing a most-general explanation for a why-not question, relative to an ontology, which can either be provided by the user, or it may be automatically derived from the data and/or schema. We study the complexity of this problem and associated problems, and present concrete algorithms for computing why-not explanations. In the case where an external ontology is provided, we first show that the problem of deciding the existence of an explanation to a why-not question is NP-complete in general. However, the problem is solvable in polynomial time for queries of bounded arity, provided that the ontology is specified in a suitable language, such as a member of the DL-Lite family of description logics, which allows for efficient concept subsumption checking. Furthermore, we show that a most-general explanation can be computed in polynomial time in this case. In addition, we propose a method for deriving a suitable (virtual) ontology from a database and/or a schema, and we present an algorithm for computing a most-general explanation to a why-not question, relative to such ontologies. This

algorithm runs in polynomial-time in the case when concepts are defined in a selection-free language, or if the underlying schema is fixed. Finally, we also study the problem of computing short most-general explanations, and we briefly discuss alternative definitions of what it means to be an explanation, and to be most general."

PODS Session 2: Awards Session

Parallel-Correctness and Transferability for Conjunctive Queries

Tom J. Ameloot, Hasselt University; Gaetano Geck, TU Dortmund University; Bas Ketsman, Hasselt University; Frank Neven, Hasselt University; Thomas Schwentick, TU Dortmund University.

A dominant cost for query evaluation in modern massively distributed systems is the number of communication rounds. For this reason, there is a growing interest in single-round multiway join algorithms where data is first reshuffled over many servers and then evaluated in a parallel but communication-free way. The reshuffling itself is specified as a distribution policy. We introduce a correctness condition, called parallel-correctness, for the evaluation of queries w.r.t. a distribution policy. We study the complexity of parallel-correctness for conjunctive queries as well as transferability of parallel-correctness between queries. We also investigate the complexity of transferability for certain families of distribution policies, including, for instance, the Hypercube distribution.

PODS Session 4: Existential Rules and TGDs

Function Symbols in Tuple-Generating Dependencies: Expressive Power and Computability

Georg Gottlob, Oxford University; Reinhard Pichler, Vienna University of Technology; Emanuel Sallinger, Vienna University of Technology.

Tuple-generating dependencies - for short tgds - have been a staple of database research throughout most of its history. Yet one of the central aspects of tgds, namely the role of existential quantifiers, has not seen much investigation so far. When studying dependencies, existential quantifiers and - in their Skolemized form - function symbols are often viewed as two ways to express the same concept. But in fact, tgds are quite restrictive in the way that functional terms can occur. In this paper, we investigate the role of function symbols in dependency formalisms that go beyond tgds. Among them is the powerful class of SO tgds and the intermediate class of nested tgds. In addition, we employ Henkin quantifiers - a well-known concept in the area of logic - and introduce Henkin tgds to gain a more fine-grained understanding of the role of function symbols in dependencies. For members of these families of dependency classes, we investigate their expressive power, that is, when one dependency class is equivalently representable in another class of dependencies. In addition, we analyze the computability of query answering under many of the well-known syntactical decidability criteria for tgds as well as the complexity of model checking.

Default Negation for Non-Guarded Existential Rules

Mario Alviano, University of Calabria; Andreas Pieris, Vienna University of Technology.

The problem of query answering under the well-founded and stable model semantics for normal existential rules, that is, existential rules enriched with default negation, has recently attracted a lot of interest from the database and KR communities. In particular, it has been thoroughly studied for classes of normal existential rules that are based on restrictions that guarantee the tree-likeness of the underlying models; a prime example of such a restriction is guardedness. However, little is known about classes of existential rules that significantly deviate from the above paradigm. A prominent example of such a formalism is the class of existential rules that is based on the notion of stickiness, which enforces restrictions on the forms of joins in the rule-bodies. It is the precise aim of the current work to extend sticky existential rules with default negation, and perform an in-depth analysis of the complexity of conjunctive query answering under the well-founded and stable model semantics. We show that an effective way for bridging the gap between stickiness and the well-founded semantics exists, and we provide data and combined complexity results. However, there is no way to reconcile stickiness and the stable model semantics. The reason for this surprising negative result should be found in the fact that sticky existential rules are powerful enough for expressing cartesian products, a construct that forms a prime example of non-guardedness.

Chase Termination for Guarded Existential Rules

Marco Calautti, University of Calabria; Georg Gottlob, Oxford University; Andreas Pieris, Vienna University of Technology.

The chase procedure is considered as one of the most fundamental algorithmic tools in database theory. It has been successfully applied to different database problems such as data exchange, and query answering and containment under constraints, to name a few. One of the central problems regarding the chase procedure is all-instance termination, that is, given a set of tuple-generating dependencies (TGDs) (a.k.a. existential rules), decide whether the chase under that set terminates, for every input database. It is well-known that this problem is undecidable, no matter which version of the chase

we consider. The crucial question that comes up is whether existing restricted classes of TGDs, proposed in different contexts such as ontological query answering, make the above problem decidable. In this work, we focus our attention on the oblivious and the semi-oblivious versions of the chase procedure, and we give a positive answer for classes of TGDs that are based on the notion of guardedness. To the best of our knowledge, this is the first work that establishes positive results about the (semi-)oblivious chase termination problem. In particular, we first concentrate on the class of linear TGDs, and we syntactically characterize, via rich- and weak-acyclicity, its fragments that guarantee the termination of the oblivious and the semi-oblivious chase, respectively. Those syntactic characterizations, apart from being interesting in their own right, allow us to pinpoint the complexity of the problem, which is PSPACE-complete in general, and NL-complete if we focus on predicates of bounded arity, for both the oblivious and the semi-oblivious chase. We then proceed with the more general classes of guarded and weakly-guarded TGDs. Although we do not provide syntactic characterizations for its relevant fragments, as for linear TGDs, we show that the problem under consideration remains decidable. In fact, we show that it is 2EXPTIME-complete in general, and EXPTIME-complete if we focus on predicates of bounded arity, for both the oblivious and the semi-oblivious chase. Finally, we investigate the expressive power of the query languages obtained from our analysis, and we show that they are equally expressive with standard database query languages. Nevertheless, we have strong indications that they are more succinct.

Recovering Exchanged Data

Gosta Grahne, Concordia University; Ali Moallemi, Concordia University; Adrian Onet, Concordia University.

The inversion of data exchange mappings is one of the thorniest issues in data exchange. In this paper we study inverse data exchange from a novel perspective. Previous work has dealt with the static problem of finding a target-to-source mapping that captures the “inverse” of a source-to-target data exchange mapping. As we will show this approach has some drawbacks when it comes actually applying the inverse mapping in order to recover a source instance from a materialized target instance. More specifically (1): As is well known, the inverse mappings have to be expressed in a much more powerful language than the mappings they invert. (2): There are simple cases where a source instance computed by the inverse mapping misses sound information that one may easily obtain when the particular target instance is available. (3): In some cases the inverse mapping can introduce unsound information in the recovered source instance. To overcome these drawbacks we focus on the dynamic problem of recovering the source instance using the source-to-target mapping as well as a given target instance. Similarly to the problem of finding “good” target instances in forward data exchange, we look for “good” source instances to restore, i.e. to materialize. For this we introduce a new semantics to capture instance based recovery. We then show that given a target instance and a source-to-target mapping expressed as set of tuple generating dependencies, there are chase-based algorithms to compute a representative finite set of source instances that can be used to get certain answers to any union of conjunctive source queries. We also show that the instance based source recovery problem unfortunately is coNP-complete. We therefore present a polynomial time algorithm that computes a “small” set of source instances that can be used to get sound certain answers to any union of conjunctive source queries. This algorithm is then extended to extract more sound information for the case when only conjunctive source queries are allowed.

PODS Session 5: Tree Structured Data

The (Almost) Complete Guide to Tree Pattern Containment

Wojciech Czerwiński, University of Warsaw; Wim Martens, University of Bayreuth; Paweł Parys, University of Warsaw; Marcin Przybytko, University of Warsaw.

Tree pattern queries are being investigated in database theory for more than a decade. They are a fundamental and flexible query mechanism and have been considered in the context of querying tree structured as well as graph structured data. We revisit their containment, validity, and satisfiability problem, both with and without schema information. We present a comprehensive overview of what is known about the complexity of containment and develop new techniques which allow us to obtain tractability- and hardness results for cases that have been open since the early work on tree pattern containment. For the tree pattern queries we consider in this paper, it is known that the containment problem does not depend on whether patterns are evaluated on trees or on graphs. This means that our results also shed new light on tree pattern queries on graphs.

Efficient Evaluation and Approximation of Well-designed Pattern Trees

Pablo Barceló, University of Chile; Reinhard Pichler, Vienna University of Technology; Sebastian Skritek, Vienna University of Technology.

Conjunctive queries (CQs) fail to provide an answer when the pattern described by the query does not exactly match the data. CQs might thus be too restrictive as a querying mechanism when data is semistructured or incomplete. The semantic web therefore provides a formalism - known as well-designed pattern trees (WDPTs) - that tackles this problem: WDPTs allow us to match patterns over the data, if available, but do not fail to give an answer otherwise. Here we abstract away the

specifics of semantic web applications and study WDPTs over arbitrary relational schemas. Our language properly subsumes the class of CQs. Hence, WDPT evaluation is intractable. We identify structural properties of WDPTs that lead to tractability of various variants of the evaluation problem. For checking if a WDPT is equivalent to one in our tractable class, we prove 2EXPTIME-membership. As a corollary, we obtain fixed-parameter tractability of (variants of) the evaluation problem. Our techniques also allow us to develop a theory of approximations for WDPTs.

BonXai: Combining the Simplicity of DTD with the Expressiveness of XML Schema

Wim Martens, University of Bayreuth; Frank Neven, Hasselt University; Matthias Niewerth, University of Bayreuth; Thomas Schwentick, TU Dortmund University.

While the migration from DTD to XML Schema was driven by a need for increased expressivity and flexibility, the latter was also significantly more complex to use and understand. Whereas DTDs are characterized by their simplicity, XML Schema Definitions (XSDs) are notoriously difficult. In this paper, we introduce the XML specification language BonXai which possesses most features of XSDs, including its expressivity, while retaining the simplicity of DTDs. In brief, the latter is achieved by sacrificing the explicit use of types in favor of simple patterns expressing contexts for elements. The goal of BonXai is by no means to replace XML Schema, but rather to provide a simpler DTD-like alternative to schema designers that do not need the explicit use of types. Therefore, BonXai can be seen as a practical front-end for XML Schema. A particular strong point of BonXai is its solid foundation rooted in a decade of theoretical work around pattern-based schemas. We present in detail the formal model for BonXai and discuss translation algorithms to and from XML Schema.

PODS Session 6: Tutorial Session 2

Defining Relations on Graphs: How Hard is it in the Presence of Node Partitions?

M Praveen, Chennai Mathematical Institute; B Srivathsan, Chennai Mathematical Institute.

Designing query languages for graph structured data is an active field of research. Evaluating a query on a graph results in a relation on the set of its nodes. In other words, a query is a mechanism for defining relations on a graph. Some relations may not be definable by any query in a given language. This leads to the following question: given a graph, a query language and a relation on the graph, does there exist a query in the language that defines the relation? This is called the definability problem. When the given query language is standard regular expressions, the definability problem is known to be PSPACE-complete. The model of graphs can be extended by labeling nodes with values from an infinite domain. These labels induce a partition on the set of nodes: two nodes are equivalent if they are labeled by the same value. Query languages can also be extended to make use of this equivalence. Two such extensions are Regular Expressions with Memory (REM) and Regular Expressions with Equality (REE). In this paper, we study the complexity of the definability problem in this extended model when the query language is either REM or REE. We show that the definability problem is EXSPACE-complete when the query language is REM, and it is PSPACE-complete when the query language is REE. In addition, when the query language is a union of conjunctive queries based on REM or REE, we show CoNP-completeness.

PODS Session 7: Query Processing

Querying Big Data by Accessing Small Data

Wenfei Fan, University of Edinburgh; Floris Geerts, University of Antwerp; Yang Cao, University of Edinburgh; Ting Deng, Beihang University; and Ping Lu, Beihang University.

This paper investigates the feasibility of querying big data by accessing a bounded amount of the data. We study boundedly evaluable queries under a form of access constraints, when their evaluation cost is determined by the queries and constraints only. While it is undecidable to determine whether FO queries are boundedly evaluable, we show that for several classes of FO queries, the bounded evaluability problem is decidable. We also provide characterization and effective syntax for their boundedly evaluable queries. When a query Q is not boundedly evaluable, we study two approaches to approximately answering Q under access constraints. (1) We search for upper and lower envelopes of Q that are boundedly evaluable and warrant a constant accuracy bound. (2) We instantiate a minimum set of variables (parameters) in Q such that the specialized query is boundedly evaluable. We study problems for deciding the existence of envelopes and bounded specialized queries, and establish their complexity for various classes of FO queries.

Skyline Queries with Noisy Comparisons

Benoit Groz, Paris-Sud University; Tova Milo, Tel-Aviv University.

We study in this paper the computation of skyline queries - a popular tool for multicriteria data analysis - in the presence of noisy input. Motivated by crowdsourcing applications, we present the first algorithms for skyline evaluation in a computation model where the input data items can only be compared through noisy comparisons. In this model comparisons

may return wrong answers with some probability, and confidence can be increased through independent repetitions of a comparison. Our goal is to minimize the number of comparisons required for computing or verifying a candidate skyline, while returning the correct answer with high probability. We design output-sensitive algorithms, namely algorithms that take advantage of the potentially small size of the skyline, and analyze the number of comparison rounds of our solutions. We also consider the problem of predicting the most likely skyline given some partial information in the form of noisy comparisons, and show that optimal prediction is computationally intractable.

The Communication Complexity of Distributed Set-Joins with Applications to Matrix Multiplication

Dirk Van Gucht, Indiana University; Ryan Williams, Stanford University; David Woodruff, IBM; Qin Zhang, Indiana University.

Given a set-comparison predicate P and given two lists of sets $A = (A_1, \dots, A_m)$ and $B = (B_1, \dots, B_m)$, with all A_i, B_j subsets of $[n]$, the P -set join $A \bowtie^P B$ is defined to be the set $\{(i, j) \text{ in } [m] \times [m] \mid P(A_i, B_j)\}$. When $P(A_i, B_j)$ is the condition “ $A_i \cap B_j$ is not empty” we call this the set-intersection-notempty join (a.k.a. the composition of A and B); when $P(A_i, B_j)$ is “ $A_i \cap B_j$ is empty” we call it the set-disjointness join; when $P(A_i, B_j)$ is “ $A_i = B_j$ ” we call it the set-equality join; when $P(A_i, B_j)$ is “ $|A_i \cap B_j| \geq T$ ” for a given threshold T , we call it the set-intersection threshold join. Assuming A and B are stored at two different sites in a distributed environment, we study the (randomized) communication complexity of computing these, and related, set-joins $A \bowtie^P B$, as well as the (randomized) communication complexity of computing the exact and approximate value of their size $k = |A \bowtie^P B|$. Combined, our analyses shed new insights into the quantitative differences between these different set-joins. Furthermore, given the close affinity of the natural join and the set-intersection-not-empty join, our results also yield communication complexity results for computing the natural join in a distributed environment. Additionally, we obtain new algorithms for computing the distributed set-intersection-not-empty join when the input and/or output is sparse. For instance, when the output is k -sparse, we improve an $O(kn)$ communication algorithm of (Williams and Yu, SODA 2014). Observing that the set-intersection-not-empty join is isomorphic to Boolean matrix multiplication (BMM), our results imply new algorithms for fundamental graph theoretic problems related to BMM. For example, we show how to compute the transitive closure of a directed graph in $O(k^{3/2})$ time, when the transitive closure contains at most k edges. When $k = O(n)$, we obtain a (practical) $O(n^{3/2})$ time algorithm, improving a recent $O(n^{1.5(\omega+3)/4})$ time algorithm (Borassi, Crescenzi, and Habib, arXiv 2014) based on (impractical) fast matrix multiplication, where $\omega \geq 2$ is the exponent for matrix multiplication.

Joins via Geometric Resolutions: Worst-case and Beyond

Mahmoud Abo Khamis, University at Buffalo; Hung Ngo, University at Buffalo; Christopher Re, Stanford University; Atri Rudra, University at Buffalo.

We present a simple geometric framework for the relational join. Using this framework, we design an algorithm that achieves the fractional hypertree-width bound, which generalizes classical and recent worst-case algorithmic results on computing joins. In addition, we use our framework and the same algorithm to show a series of what are colloquially known as beyond worst-case results. The framework allows us to prove results for data stored in Btrees, multidimensional data structures, and even multiple indices per table. A key idea in our framework is formalizing the inference one does with an index as a type of geometric resolution; transforming the algorithmic problem of computing joins to a geometric problem. Our notion of geometric resolution can be viewed as a geometric analog of logical resolution. In addition to the geometry and logic connections, our algorithm can also be thought of as backtracking search with memoization.

PODS Session 8: Streaming

External Memory Stream Sampling

Xiaocheng Hu, the Chinese University of Hong Kong; Miao Qiao, the Chinese University of Hong Kong; Yufei Tao, the Chinese University of Hong Kong.

This paper aims to understand the I/O-complexity of maintaining a *big sample set* - whose size exceeds the internal memory's capacity - on a data stream. We study this topic in a new computation model, named the *external memory stream (EMS) model*, that naturally extends the standard external memory model to stream environments. A suite of EMS-indigenous techniques are presented to prove matching lower and upper bounds for with-replacement (WR) and without-replacement (WoR) sampling on append-only and time-based sliding window streams, respectively. Our results imply that, compared to RAM, the EMS model is perhaps a more suitable computation model for studying stream sampling, because the new model separates different problems by their hardness in ways that could not be observed in RAM.

Vertex and Hyperedge Connectivity in Dynamic Graph Streams

Sudipto Guha, University of Pennsylvania; Andrew McGregor, University of Massachusetts; David Tench, University of Massachusetts.

A growing body of work addresses the challenge of processing dynamic graph streams: a graph is defined by a sequence of edge insertions and deletions and the goal is to construct synopses and compute properties of the graph while using only

limited memory. Linear sketches have proved to be a powerful technique in this model and can also be used to minimize communication in distributed graph processing. We present the first linear sketches for estimating vertex connectivity and constructing hypergraph sparsifiers. Vertex connectivity exhibits markedly different combinatorial structure than edge connectivity and appears to be harder to estimate in the dynamic graph stream model. Our hypergraph result generalizes the work of Ahn et al. (PODS 2012) on graph sparsification and has the added benefit of significantly simplifying the previous results. One of the main ideas is related to the problem of reconstructing subgraphs that satisfy a specific sparsity property. We introduce a more general notion of graph degeneracy and extend the graph reconstruction result of Becker et al. (IPDPS 2011).

Fast and Near-Optimal Algorithms for Approximating Distributions by Histograms

Jayadev Acharya, MIT; Ilias Diakonikolas, University of Edinburgh; Chinmay Hegde, MIT; Jerry Li, MIT; Ludwig Schmidt, MIT.

Histograms are among the most popular structures for the succinct summarization of data in a variety of database applications. In this work, we provide fast and near-optimal algorithms for approximating arbitrary one dimensional data distributions by histograms. A k -histogram is a piecewise constant function with k pieces. We consider the following natural problem, previously studied by Indyk, Levi, and Rubinfeld in PODS 2012: given samples from a distribution p over $\{1, \dots, n\}$, compute a k histogram that minimizes the l_2 -distance from p , up to an additive ϵ . We design an algorithm for this problem that uses the information-theoretically minimal sample size of $m = O(1/\epsilon^2)$, runs in sample-linear time $O(m)$, and outputs an $O(k)$ -histogram whose l_2 -distance from p is at most $O(\text{opt}_k) + \epsilon$, where opt_k is the minimum l_2 -distance between p and any k -histogram. Perhaps surprisingly, the sample size and running time of our algorithm are independent of the universe size. We generalize our approach to obtain fast algorithms for multi-scale histogram construction, as well as approximation by piecewise polynomial distributions. We experimentally demonstrate one to two orders of magnitude improvement in terms of empirical running times over previous state-of-the-art algorithms.

PODS Session 9: Algorithms and Data Structures

On Top- k Range Reporting in 2D Space

Rahul Saladi, University of Minnesota; Yufei Tao, Chinese University of Hong Kong.

Orthogonal range reporting (ORR) is a classic problem in computational geometry and databases, where the objective is to preprocess a set P of points in \mathbb{R}^2 such that, given an axis-parallel rectangle q , all the points in $P \cap q$ can be reported efficiently. This paper studies a natural variant of the problem called *top- k ORR*, where each point $p \in P$ carries a *weight* $w(p) \in \mathbb{R}$. Besides q , a query also specifies an integer $k \in [1, |P|]$, and needs to report the k points in $q \cap P$ with the largest weights. We present optimal or near-optimal structures for solving the top- k ORR problem in the pointer machine and external memory models. As a side product, our structures give new space-query tradeoff for the *orthogonal range max* problem, which is a special case of top- k ORR with $k = 1$.

Dynamic Data Structures for Document Collections and Graphs

Ian Munro, University of Waterloo; Yakov Nekrich, University of Waterloo; Jeff Vitter, University of Kansas.

In the dynamic indexing problem, we must maintain a changing collection of text documents so that we can efficiently support insertions, deletions, and pattern matching queries. We are especially interested in developing efficient data structures that store and query the documents in compressed form. All previous compressed solutions to this problem rely on answering rank and select queries on a dynamic sequence of symbols. Because of the lower bound in [Fredman and Saks, 1989], answering rank queries presents a bottleneck in compressed dynamic indexing. In this paper we show how this lower bound can be circumvented using our new framework. We demonstrate that the gap between static and dynamic variants of the indexing problem can be almost closed. Our method is based on a novel framework for adding dynamism to static compressed data structures. Our framework also applies more generally to dynamizing other problems. We show, for example, how our framework can be applied to develop compressed representations of dynamic graphs and binary relations.

Join Dependency Testing, Loomis-Whitney Join, and Triangle Enumeration

Xiaocheng Hu, Chinese University of Hong Kong; Miao Qiao, Chinese University of Hong Kong; Yufei Tao, Chinese University of Hong Kong.

In this paper, we revisit two fundamental problems in database theory. The first one is called *join dependency (JD) testing*, where we are given a relation r and a JD , and need to determine whether the JD holds on r . The second problem is called *JD existence testing*, where we need to determine if there exists *any* non-trivial JD that holds on r . We prove that JD testing is NP-hard even if the JD is defined *only* on binary relations (i.e., each with only two attributes). Unless $P = NP$, this result puts a negative answer to the question whether it is possible to efficiently test JD s defined exclusively on *small* (in terms of attribute number) relations. The question has been open since the classic NP-hard proof of Maier, Sagiv, and Yannakakis in JACM'81 which requires the JD to involve a relation of $\Omega(d)$ attributes, where d is the number of attributes in r . For JD

existence testing, the challenge is to minimize the computation cost because the problem is known to be solvable in polynomial time. We present a new algorithm for solving the problem I/O-efficiently in the external memory model. Our algorithm in fact settles the closely related *Loomis-Whitney (LW) enumeration problem*, and as a side product, achieves the optimal I/O complexity for the *triangle enumeration problem*, improving a recent result of Pagh and Silvestri in PODS'14.

PODS Session 10: Weights and Similarity

Private Multiplicative Weights Beyond Linear Queries

Jonathan Ullman, Columbia University.

A wide variety of fundamental data analyses in machine learning, such as linear and logistic regression, require minimizing a convex function defined by the data. Since the data may contain sensitive information about individuals, and these analyses can leak that sensitive information, it is important to be able to solve convex minimization in a privacy-preserving way. A series of recent results show how to accurately solve a single convex minimization problem in a differentially private manner. However, the same data is often analyzed repeatedly, and little is known about solving multiple convex minimization problems with differential privacy. For simpler data analyses, such as linear queries, there are remarkable differentially private algorithms such as the private multiplicative weights mechanism (Hardt and Rothblum, FOCS 2010) that accurately answer exponentially many distinct queries. In this work, we extend these results to the case of convex minimization and show how to give accurate and differentially private solutions to exponentially many convex minimization problems on a sensitive dataset.

Symmetric Weighted First-Order Model Counting

Paul Beame, University of Washington; Guy Van den Broeck, KU Leuven; Eric Gribkoff, University of Washington; Dan Suciu, University of Washington.

The FO Model Counting problem (FOMC) is the following: given a sentence Φ in FO and a number n , compute the number of models of Φ over a domain of size n ; the Weighted variant (WFOMC) generalizes the problem by associating a weight to each tuple and defining the weight of a model to be the product of weights of its tuples. In this paper we study the complexity of the symmetric WFOMC, where all tuples of a given relation have the same weight. Our motivation comes from an important application, inference in Knowledge Bases with soft constraints, like Markov Logic Networks, but the problem is also of independent theoretical interest. We study both the data complexity, and the combined complexity of FOMC and WFOMC. For the data complexity we prove the existence of an FO^2 formula for which FOMC is $\#P_1$ -complete, and the existence of a Conjunctive Query for which WFOMC is $\#P_1$ -complete. We also prove that all γ -acyclic queries have polynomial time data complexity. For the combined complexity, we prove that, for every fragment FO^k , $k \geq 2$, the combined complexity of FOMC (or WFOMC) is $\#P$ -complete.

Smooth Tradeoffs between Insert and Query Complexity in Nearest Neighbor Search

Michael Kapralov, IBM Watson.

Locality Sensitive Hashing (LSH) has emerged as the method of choice for high dimensional similarity search, a classical problem of interest in numerous applications. LSH-based solutions require that each data point be inserted into a number A of hash tables, after which a query can be answered by performing B lookups. The original LSH solution of [IM98] showed for the first time that both A and B can be made sublinear in the number of data points. Unfortunately, the classical LSH solution does not provide any tradeoff between insert and query complexity, whereas for data (respectively, query) intensive applications one would like to minimize insert time by choosing a smaller A (respectively, minimize query time by choosing a smaller B). A partial remedy for this is provided by Entropy LSH [Pan06], which allows to make either inserts or queries essentially constant time at the expense of a loss in the other parameter, but no algorithm that achieves a smooth tradeoff is known. In this paper, we present an algorithm for performing similarity search under the Euclidean metric that resolves the problem above. Our solution is inspired by Entropy LSH, but uses a very different analysis to achieve a smooth tradeoff between insert and query complexity. Our results improve upon or match, up to lower order terms in the exponent, best known data-oblivious algorithms for the Euclidean metric."

SIGMOD TRAVEL AWARDS

PhD Symposium

David Broneske	OvGU FIN ITI, Germany
Antonio Maccioni	University Roma Tre, Italy
Xiu Susie Fang	The University of Adelaide, Australia
Yuhong Li	University of Macau, China
Qi Fan	National University of Singapore, Singapore
Antoine Amarilli	Institut Télécom; Télécom ParisTech; CNRS LTCI, France
Tianhong Song	University of California Davis, USA
Sebastian Dorok	Otto von Guericke University Magdeburg, Germany
Kwan Hui Lim	The University of Melbourne, Australia

Undergraduate Poster

Oreoluwa Babarinsa	Harvard University, USA
Trevor Clinkenbeard	University of Waterloo, Canada
Styliani Pantela	Harvard University, USA
Adam Perelman	Stanford University, USA
Susan Tu	Stanford University, USA
Mansheng Yang	National University of Singapore, Singapore

Programming Contest

Alexey Karyakin	University of Waterloo, Canada
Ingo Müller	Karlsruhe Institute of Technology, Germany
Eunjin Song	Seoul National University, South Korea
Thai Son Nguyen	Karlsruhe Institute of Technology, Germany
Lambros Petrou	University of Oxford, UK

Travel Grant

Xiaolan Wang	University of Massachusetts, Amherst, USA
Andrea Lottarini	Columbia University, USA
Natali Ruchansky	Boston University, USA
Fotios Psallidas	Columbia University, USA
Jingbo Shang	University of Illinois, Urbana-Champaign, USA
Cory Thoma	University of Pittsburgh, USA
Faisal Nawab	University of California Santa Barbara, USA
Zhiting Hu	Carnegie Mellon University, USA
Dana Van Aken	Carnegie Mellon University, USA
Shumo Chu	University of Washington, USA
Prasang Upadhyaya	University of Washington, USA

Leilani Battle	MIT, USA
Kerim Yasin Oktay	University of California Irvine, USA
Orestis Polychroniou	Columbia University, USA
Nick R. Katsipoulakis	University of Pittsburgh, USA
Vaibhav Arora	University of California Santa Barbara, USA
John Paparrizos	Columbia University, USA
Charalampos Mavroforakis	Boston University, USA
Vera Moffitt	Drexel University, USA
Shi Qiao	Case Western Reserve University, USA
Isabelle Hang	University of Bremen, Germany
Michael Cochez	University of Jyväskylä, Finland
Tobias Mühlbauer	Technical University of Munich, Germany
Xu Chu	University of Waterloo, Canada
Ahmed El-Roby	University of Waterloo, Canada
Claude Bartheles	ETH Zurich, Switzerland
Jules Chevalier	Laboratoire Hubert Curien, France
Jian He	Tsinghua University, China
Xiaoyang Wang	University of New South Wales, Australia
Fangzhou Zhu	Soochow University, China
Yue Wang	Chinese Academy of Sciences, China
Yudian Zheng	The University of Hong Kong, China
Minhao Jiang	Hong Kong University of Science and Technology, China
Yanxiang Huang	Peking University, China
Jun Zhang	Nanyang Technological University, Singapore
Ziqiang Feng	The Hong Kong Polytechnic University, China
Kijung Shin	Seoul National University, South Korea
Saket Gururkar	Indian Institute of Technology Madras, India
Furong Li	National University of Singapore, Singapore
Wen Hua	Renmin University of China, China
Peng Peng	Hong Kong University of Science and Technology, China
Guo Tao	Nanyang Technological University, Singapore
Qian Lin	National University of Singapore, Singapore
Ruifeng Liu	The Chinese University of Hong Kong, China
Jinhong Jung	KAIST (Korea Advanced Institute of Science and Technology), South Korea

PH.D. WORKSHOP PARTICIPANTS

David Broneske (OvGU FIN ITI, Germany)

Antonio Maccioni (University Roma Tre, Italy)

Xiu Susie Fang (The University of Adelaide, Australia)

Yuhong Li (University of Macau, China)

Qi Fan (National University of Singapore, Singapore)

Antoine Amarilli (Institut Télécom; Télécom ParisTech; CNRS LTCI, France)

Tianhong Song (University of California Davis, USA)

Sebastian Dorok (Otto von Guericke University Magdeburg, Germany)

Kwan Hui Lim (The University of Melbourne, Australia)

CO-LOCATED WORKSHOPS

SIGMOD/PODS Ph.D. Symposium

<http://staff.itee.uq.edu.au/shenht/phdsymposium2015/>

May 31, 2015

Location: MR106

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

Workshop on Data analytics at sCale (DanaC)

<http://danac.org/>

May 31, 2015

Location: MR104

Data nowadays comes from various sources including log files, transactional applications, the Web, social media, scientific experiments, and many others. Much of this data is generated and transmitted in real time and at an unprecedented scale. To extract value out of these data sets, business analysts and data scientists employ advanced data analytics techniques combining, among others, traditional BI, text analytics, machine learning, data mining and many others. However, we are still far from being able to ingest, query and visualize large amounts of data in a scalable, easy to use and standardized manner. To answer these challenges, every day, new, more complex analysis techniques are devised and larger, more varied datasets are accumulated. Tackling the complexity of both the data itself and its analysis remains an open challenge. The target audience of the workshop is database, storage, networking, and distributed systems researchers, domain scientists using Big Data for their research, practitioners building systems for Big Data processing, as well as data analysts that extract value out of Big Data.

Workshop on Data Management on New Hardware (DaMoN)

<http://event.cwi.nl/damon2015/>

June 1, 2015

Location: MR103

The aim of this one-day workshop is to bring together researchers who are interested in optimizing database performance on modern computing infrastructure by designing new data management techniques and tools.

Workshop on Exploratory Search in Databases and the Web (ExploreDB)

<https://sites.google.com/site/exploredb2015/>

May 31, 2015

Location: MR107

Nowadays, there is a need to develop novel paradigms for exploratory user-data interactions that emphasize user context and interactivity with the goal of facilitating exploration, interpretation, retrieval, and assimilation of information. A huge number of applications need an exploratory form of querying. Ranked retrieval techniques for relational databases, XML, RDF and graph databases, text and multimedia databases, scientific and statistical databases, social networks and many others, is a first step towards this direction. Recently, several new aspects for exploratory search, such as preferences, diversity, novelty, surprise and serendipity, are gaining increasing importance. From a different perspective, recommendation applications tend to anticipate user needs by automatically suggesting the information which is most appropriate to the users and their current context. Also, a new line of research in the area of exploratory search is fueled by the growth of online social interactions within social networks and Web communities. Many useful facts about entities (e.g., people, locations, organizations, products) and their relationships can be found in a multitude

of semi-structured and structured data sources, such as Wikipedia, Linked Open Data cloud, Freebase, and many others. Therefore, novel discovery methods are required to provide highly expressive discovery capabilities over large amounts of entity-relationship data, which are yet intuitive for end-users.

Workshop on Managing and Mining Enriched Geo-Spatial Data (GeoRich)

<http://www.dbs.ifi.lmu.de/georich15/>

May 31, 2015

Location: MR101&102

Both the current trends in technology, such as smart phones, general mobile devices, stationary sensors and satellites, as well as a new user mentality of utilizing this technology to voluntarily share information, cause the production of a huge flood of geo-spatial data. This data can integrate multiple other types of data, e.g., social information, textual data, multimedia data and scientific measurements, leading to what one can call generally as "enriched geospatial data". This flood of geo-referenced data provides a tremendous potential of discovering new and possibly useful knowledge. A novel research challenge is thus to search and mine this wealth of enriched geo-spatial data which requires joint effort from multiple research communities.

Graph Data-management Experiences & Systems (GRADES)

<http://event.cwi.nl/grades2015/>

May 31, 2015

Location: MR103

The focus of this new workshop is to elicit discussion and description of application areas and open challenges in managing large-scale graph-shaped data.

Experiences including use-case descriptions, war stories and benchmarks as well as system descriptions are important ingredients of the meeting.

Workshop on the Web and Databases (WebDB)

<http://dbweb.enst.fr/events/webdb2015/index.html>

May 31, 2015

Location: MR105

The Web is changing during the last decade, and becoming an evolving source of information and knowledge, with content being derived, generated, shared or simply annotated by end-users. The heterogeneity of data formats and applications affects the way knowledge is produced and consumed. This shift from the old Web, in which knowledge was statically generated, to the new Web, where semantics can be exploited by ordinary users, who actively participate in knowledge creation, dissemination and maintenance, is presenting new opportunities and challenges. The focus of this year's WebDB workshop is on novel techniques for modeling and reasoning about freshness, correctness and quality of information and knowledge on the Web. We solicit contributions from the data management community, in which traditional data management techniques are extended, adapted and extensively tested, in support of a transition from the Web of Data to the Web of Knowledge.

LOCAL INFORMATION

Melbourne is an exciting, ever-changing city with extraordinary surprises to be discovered in every laneway, rooftop and basement.

A city that regularly tops the World's Most Liveable list, there's always something going on with an enviable array of major events and festivals.

Set on the banks of the iconic Yarra River, Melbourne Convention and Exhibition Centre (MCEC) is just a 20 minute drive from Melbourne airport and a short stroll to the city centre. There you can explore cafés, eclectic restaurants, lush parklands, incredible fashion and inspiring galleries and museums.

Things to do in Melbourne:

- Melbourne Visitor Centre located at Federation Square is a great starting point for exploring the city
- Mission to Seafarers – Completed in 1919, the heritage-listed Mission to Seafarers Victoria building is a fine example of Arts and Crafts, Spanish Mission Revival architecture (717 Flinders Street, Docklands)
- Polly Woodside – a historic museum ship moored at South Wharf, a tangible reminder of Australia's rich maritime heritage (right outside the Convention Centre)
- Queen Victoria Market – more than the city's fresh-food shopping mecca – it's a historic landmark, a tourist attraction and a Melbourne institution (corner Victoria St and Elizabeth St, Melbourne)
- The Royal Arcade has provided an undercover shopping link between Melbourne's Bourke Street Mall, Little Collins Street and Elizabeth Street since 1869.
- Shrine of Remembrance – the Shrine is a Melbourne icon. Guided tours run from the visitor centre daily at 11am and 2pm, and Services of Remembrance are held every half-hour beginning at 10.30am (Birdwood Ave, Melbourne)
- Eureka Skydeck 88 – visit the Southern Hemisphere's highest viewing platform at Eureka Skydeck 88, with 360-degree views. Ride the lift to level 88 in under 40 seconds, then step onto the observation deck for amazing floor-to-ceiling views that take in the whole of Melbourne, from Docklands and Port Phillip Bay to the Dandenong Ranges and Mornington Peninsula.
- Melbourne Star – Australia's only giant observation wheel. Take a flight of discovery through the Melbourne skies and see the city from a different perspective (101 Waterfront Way, Docklands)
- NGV – The National Gallery of Victoria (NGV) is the oldest and most visited gallery in Australia. Situated over two magnificent buildings – NGV International and NGV Australia – the Gallery hosts a wide range of international and local artists, exhibitions, programs and events; from contemporary art to major international historic exhibitions, fashion and design, architecture, sound and dance (180 St Kilda Road, Melbourne)
- ACMI – A unique institution at the heart of Melbourne's iconic meeting place, Federation Square, the Australian Centre for the Moving Image (ACMI) celebrates, explores and promotes the moving image in all its forms - film, television and digital culture.

Restaurants Close to MCEC

Breakfast & Lunch

Caffe Cino

With a menu that focuses on fresh, local produce prepared with striking simplicity, Caffe Cino at Hilton Melbourne South Wharf is the perfect light-filled, casual space for a quick bite or relaxing meeting.
Part 2 Convention Centre Place SOUTH WHARF

Charlie Lovett

Charlie Lovett is a coffee roaster, baker and pizza maker. Enjoy store roasted, award-winning coffee beans or a piping-hot Belgian hot chocolate.
43 South Wharf Promenade SOUTH WHARF

Citizen

The understated Citizen cafe and bar, on the city's edge, serves chic cafe fare and its signature blend of coffee, enjoyed indoors or alfresco. Sit back at this cafe within Melbourne Public, at South Wharf Promenade, and watch as the city backdrop unfolds and the promenade comes to life.

The Boatbuilders Yard

The Boatbuilders Yard bar and cafe brings relaxed and informal eating and drinking to Melbourne, with plenty of European style and finesse.
Part 23 South Wharf Promenade SOUTH WHARF

Lunch & Dinner

Akachochin

Named after the pretty red paper lanterns that adorn Japanese eateries, Akachochin is an izakaya-style restaurant specialising in everything Japanese.
33 Dukes Walk SOUTHBANK

BangPop

BangPop brings the energetic vibe of Bangkok's hawker street stalls to South Wharf, emphasising communal dining, family recipes and spicy flavours.
35 South Wharf Promenade SOUTH WHARF

Baraonda

Baraonda takes pride in the quality and freshness of its Italian-style food.
18 Siddeley Street DOCKLANDS

Bohemian Bar & Restaurant

The Bohemian is a Spanish-influenced brasserie – a little shot of Madrid to stimulate the senses with delicious fine food and authentic surrounds. The tapas, raciones and postres will keep you coming back for more.
35 Dukes Walk SOUTH WHARF

Byblos Bar and Restaurant

Byblos Bar and Restaurant combines imaginative Mediterranean cuisine with lavish hospitality in luxurious surroundings overlooking the Yarra River.
18-38 Siddeley St MELBOURNE

Gasolina

Combining quality cuisine with an innate passion for motorcycles, Gasolina is a culture-rich restaurant and bar.
25 Dukes Walk SOUTH WHARF

Meat Market South Wharf

A unique experience inspired by the passion of the Aussie BBQ, Meat Market at South Wharf has an open kitchen at the centre of all the action.

53 South Wharf Promenade SOUTH WHARF

Sotano Wine and Tapas

Sotano Wine and Tapas in the Hilton at South Wharf showcases authentic tapas bites, complemented by a creative range of Spanish-inspired cocktails.

The General Assembly

The philosophy at The General Assembly is to keep things simple, with high quality, well-sourced products; and don't mess with what's already a good thing.

29 South Wharf Promenade SOUTH WHARF

Bars**Atrium Bar**

The elegant Atrium Bar at Crown is an indulgent champagne bar with a sophisticated array of premium champagnes on offer.

Part Casino 8 Whiteman Street SOUTHBANK

Lumia Bar

Lumia is a lavish cocktail bar – a place for being seen in all your glamour. Marvel at breathtaking views of the city skyline and river, from the outdoor terrace.

The West End, Crown Casino, 8 Whiteman Street SOUTHBANK

Melbourne Public

Melbourne Public celebrates the best of our past in true pub style.

11 Dukes Wharf SOUTH WHARF

Munich Brauhaus

Munich Brauhaus is a 900-seater German beerhall in the heart of South Wharf. Enjoy Oktoberfest in Melbourne, all year round.

45 South Wharf Promenade SOUTH WHARF

Plus 5 Bar and Restaurant

The cocktail menu features both the classics and a host of unique signature cocktails. The food is a mix of charcuterie and small bites, through to a tight pizza list.

37 South Wharf Promenade SOUTH WHARF

What's on Melbourne

Unusual Wine Variety Dinner

12 Jun 2015, The Wayside Inn

Top Australian Wines and The Wayside Inn present a food and wine matching event that focuses on unusual, flavour-forward pairings

Battlefield Gallipoli: Landscape of War and Memory

14 Apr 2015 - 31 Aug 2015, Shrine of Remembrance

The enduring fascination with Gallipoli is explored in this exhibition which presents some 100 items recovered from an archaeological survey project.

Anything Goes

27 May 2015 - 5 Jul 2015, Princess Theatre

Helpmann Award-winning Dean Bryan and Andrew Hallsworth's brand-new production of Cole Porter's delightful, de-lovely musical, Anything Goes.

Shifting Gear: Design, Innovation and the Australian Car

6 Mar 2015 - 12 Jul 2015, Ian Potter Centre: NGV Australia

An exhibition revealing iconic Australian-designed vehicles as objects of art and design.

Urban Koorie Tour

10 Apr 2015 - 26 Jun 2015, Federation Square

Learn how Aboriginal people used the land that we now call Melbourne in a guided tour.

Good Food and Wine Show

5 Jun 2015 - 8 Jun 2015, Melbourne Convention and Exhibition Centre

Join hundreds of local and international exhibitors as they present the best in food and wine Australia has to offer.

A Golden Age of China: Qianlong Emperor 1736-1795

27 Mar 2015 - 21 Jun 2015, NGV International

Explore the hidden treasures of the Beijing Palace Museum art collection at this exclusive exhibition.

Strictly Ballroom the Musical

20 Jan 2015 - 26 Jul 2015, Her Majesty's Theatre

The brand-new production of Strictly Ballroom the Musical has audiences dancing in the aisles in a kaleidoscope of colour and fun.

Melbourne International Jazz Festival

28 May 2015 - 7 Jun 2015, Arts Centre Melbourne

The 2015 Melbourne International Jazz Festival (MIJF) presents an incredible program featuring some of the world's leading jazz artists.

First Peoples

7 Sep 2013 - 31 Dec 2023, Melbourne Museum

First Peoples celebrates the history, culture, achievements and survival of Victoria's Aboriginal people.

Melbourne Street Art Tours

24 Jan 2014 - 18 Jun 2016, Federation Square

Take an artful walk with Melbourne Street Art Tours, run by street artists and providing an overview of the Melbourne underground street art scene.



This cloud stands up
to any storm.

Microsoft Azure scales to enable AccuWeather to respond to 9.5 billion requests for crucial weather data per day. This cloud rises to the challenge when the weather is at its worst.


This is the Microsoft Cloud.

learn more at microsoftcloud.com

 Microsoft Cloud

Push a Button

Move Your Database to the Oracle Cloud



... or Back to Your Data Center

ORACLE®

cloud.oracle.com/database
or call 1.800.ORACLE.1

Copyright © 2015, Oracle and/or its affiliates. All rights reserved. Oracle and Java are registered trademarks of Oracle and/or its affiliates.

M_415M_CLD00282_PushButton_DB

Research at Facebook

Giving people the power to share and connect requires constant innovation. At Facebook, research permeates everything we do. Research is more than a lab - it's a way of doing things.

We believe the most interesting academic problems are derived from real-world problems. Our engineers work on cutting edge research problems with a practical focus and push product boundaries every day. At the same time, they are publishing papers, giving talks, attending and hosting conferences, and collaborating with the academic community.



We are Hiring!

View our open positions at:
facebook.com/careers/teams/engineering

facebook

Our engineers work to solve real-world problems that impact millions of people in areas such as:

- Compilers
- Computer Vision
- Data Mining
- Databases
- Distributed Systems
- Human-Computer Interaction
- Internet Economics
- Machine Learning
- Natural Language Processing
- Networking
- Operating Systems
- Programming Languages
- Security/Privacy
- Social Computing
- Artificial Intelligence



Google's Mission To organize the world's information and make it universally accessible and useful.

Google is full of smart people working on some of the most difficult problems in Computer Science today. Many of our engineers have PhDs and come from a research background - but we don't see the theoretical and practical as being mutually exclusive. That's why research happens across the company, on many different teams, making the line between research and development at Google wonderfully blurred.

Most people know about the research activities that back our major products, such as search algorithms, systems infrastructure, machine learning, and programming languages. Those are just the tip of the iceberg; Google has a tremendous number of exciting challenges that only arise through the vast amount of data and sheer scale of systems we build. What we discover affects the world both through better Google products and services, and through dissemination of our findings by the broader academic research community. We value each kind of impact, and often the most successful projects achieve both.

We embrace our mission by hiring smart people to help us organize the world's information to make it universally accessible and useful.

Google Inc. research.google.com

**90% OF THE WORLD'S
DATA WAS CREATED IN
THE LAST TWO YEARS.
ARE YOU READY
TO HARNESS IT?**

*"Big Data, for better or worse: 90% of data
was generated over the last two years
ScienceDaily
May 22nd, 2013.*

IBM

#makedatawork

In-Memory Database Research Leader

SAP database and technology portfolio opens a completely new, exciting research field and provides the best opportunity to attract the world's attention. Join us to help the world run simpler and improve people's lives.

Find out about SAP In-Memory Database Solutions at saphana.com.

Get started at sap.com/careers

SAP



Helping people see and understand their data



tableau.com



There is tons of great content on Twitter for everyone, but at times it can be challenging to discover that content. At the same time, the opportunity is huge: we have large amounts of data from Tweet engagements both on and off Twitter which is unused today. We can use this data to improve the product experience, and to further fulfill our core mission of instantly connecting users to the information most meaningful to them without barriers.

We are looking for talented system engineers, database kernel engineers, and data scientists to come help us continue to work on solving this challenging problem.

Data scientists will build engagement prediction models for Tweets using the troves of engagement data we have, and doing ad-hoc analysis of data to generate meaningful product insights. Through this work, you will have the opportunity to meaningfully move Twitter's core product metrics


System engineers and database engineers will work on building new large scale systems capable of processing several hundred petabytes of data spread across multiple data centers for real time and batch analytics e.g. Heron. You will get an opportunity to work with talented engineers on the cutting edge open source data processing systems e.g., Hadoop and Spark, and also get to open-source a lot of your work. Twitter is one of the leaders in this space as evidenced by our open-source creations such as Scalding, Summingbird, and Parquet.

Interested in learning more? Email jointheflock@twitter.com!

Work life that benefits life life.
#yourewhatnext




Intel Labs is Looking for Research Scientists!
Submit your resume to: intellabsjobs@intel.com

Computer Science and IT

Take IT to the Limit



www.rmit.edu.au



Big Data Research




ELSEVIER

FOR MORE INFORMATION ON HOW TO SUBMIT AND VIEW THE SAMPLE ISSUE, VISIT: www.elsevier.com/locate/bdr

SIGMOD/PODS 2015 ON SOCIAL MEDIA

Please use the hashtags `#sigmod15` and/or `#pods15` 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 `#sigmod15` and `#pods15` will uniquely identify a paper:

SIGMOD Research Track Papers	<code>#sigmod15 #RxPy</code> where x is the session number (1–24) and y is the paper number within the session
SIGMOD Industrial Track Papers	<code>#sigmod15 #IxPy</code> where x is the session number (1–5) and y is the paper number within the session
SIGMOD Demos	<code>#sigmod15 #demoXY</code> where X is the session letter (A–C), and Y is the demo number within the session
SIGMOD Tutorials	<code>#sigmod15 #tutorialX</code> where X is the tutorial number (1–4)
SIGMOD Keynotes	<code>#sigmod15 #keynoteX</code> where X is the keynote number (1–2)
SIGMOD Panel	<code>#sigmod15 #panel</code>
SIGMOD New Researcher Symposium	<code>#sigmod15 #nrs</code>
SIGMOD Session on Careers in Industry	<code>#sigmod15 #career</code>
SIGMOD Award Talks Session	<code>#sigmod15 #award</code>
ACM-W Athena Lecturer Award	<code>#sigmod15 #athenaaward</code>
SIGMOD Programming Contest	<code>#sigmod15 #pc</code>
PODS Papers and Tutorials	<code>#pods15 #RxPy</code> where x is the session number (1–10) and y is the number of the presentation within the session
PODS Keynote	<code>#pods15 #keynote</code>

SIGMOD facebook page: <https://www.facebook.com/SIGMOD2015>

QR Code for SIGMOD website: <http://www.sigmod2015.org>



QR Code for SIGMOD facebook page: <https://www.facebook.com/SIGMOD2015>



SIGMOD



San Francisco 2016



SEE YOU NEXT YEAR

June 26 – July 1, 2016

San Francisco, USA

<http://www.sigmod2016.org/>