VLDB 2011

37th International Conference on Very Large Data Bases Aug 29–Sept 3: The Westin, Seattle, WA



TABLE OF CONTENTS

| Welcome Message form General Chair 1 |
|---|
| Welcome Message from the Program Chairs 2-3 |
| General Information |
| Social Events |
| Conference Venue |
| Keynotes 8-9 |
| 10 Year Best Paper Award10-11 |
| Program at a Glance |
| Conference Program |
| Tutorials |
| Panels |
| Demonstrations |
| Abstracts |
| Research Session Abstracts48-84 |
| Industrial Session Abstracts |
| Challenges & Vision Abstracts |
| Demonstration Abstracts |
| PhD Workshop Abstracts 101-103 |
| Co-Located Workshops |
| ТРСТС 104-105 |
| QBD |
| DMSN |
| DBRank/MUD |
| DBPL |
| ADMS |
| |
| DIR TE |
| VI DS 166-117 |
| PersDB |
| Conference Organization |
| Seattle Information |
| Sponsors |

WELCOME MESSAGE FROM THE GENERAL CHAIR

On behalf of the conference committee for VLDB 2011, it is my pleasure to welcome you to Seattle and the state of Washington for the 37th International Conference on Very Large Data Bases.

VLDB is a premier international forum for data management and database researchers, vendors, practitioners, application developers, and users. The conference will feature two keynotes, research talks, industry presentations, panels, tutorials, demonstrations, and eleven workshops. The program covers a wide range of topics in data management, database and information systems research. Data management and databases remain one of the key technical underpinnings of a broad variety of applications.

Seattle is the largest city in Washington State with over 3.3 million residents in the Seattle metropolitan area, around half of the state's population. The city is situated on the shores of Puget Sound, a long narrow inlet from the Pacific Ocean. It has a vibrant economy with many well-known companies in aerospace, software, and retailing. University of Washington is one of the top research universities in the country and is particularly renowned for its medical education, care and research. Many companies and research centers in the medical and biological fields are clustered near the university. Washington State is also a major fruit and wine producer with vast orchards and vineyards east of the Cascade Mountains. Don't miss the opportunity to sample some excellent Washington wines!

The conference would not have been possible without the efforts of many people. Thanks are due to the Technical Program Chairs - Dr. Jose Blakeley and Dr. Joe Hellerstein, and their team of track program chairs: Research track – Dr. Nick Koudas, Dr. Wolfgang Lehner, and Dr. Sunita Sarawagi; Industrial and Applications track – Dr. Berthold Reinwald and Dr. Phil Bohannon; Demo track – Dr. Jignesh Patel and Dr. Masatoshi Yoshikawa; and Tutorial track – Dr. Qiong Lou, and Dr. Gerome Miklau. Thanks also for the efforts of the workshop co-chairs Dr. Luis Gravano and Dr. Chris Jermaine and the Ph.D. Workshop chair Dr. Philippe Bonnet. Many, many thanks go to Dr. Uwe Roehm (Proceedings chair), Dr. Arnd Christian König (Local Organization chair), Dr. Magda Balazinska (Treasurer), Dr. David Dewitt (Sponsorship chair), Dr. Justin Levandoski (Webmaster), Dr. Yasin Silva (Program Booklet organizer), and the student volunteers. The contributions of the University of Washington Conference Management Team, ably led by Ms. Jan Kvamme, were invaluable and saved the organizers many sleepless nights.

We are grateful for the generous support of our sponsors: Platinum Sponsors – Facebook, Google, Microsoft and Oracle; Gold Sponsors - EMC, IBM, SAP, and Sybase; Silver Sponsors - HP, NEC, TPC, and Yahoo!, and Workshop Sponsors – Intelius (QDB) and Microsoft (LADIS).

Last but not least, our sincere thanks go to the authors of papers and demos, the many reviewers, the speakers, and everybody who contributed to the program. Without your contributions there would be no conference.

1

Welcome and enjoy the conference and your time in Seattle!

Per-Åke (Paul) Larson VLDB 2011 General Chair

Message from Program Chairs

Welcome to VLDB 2011 in Seattle! After a year of hard work from authors, reviewers, conference and journal organizers, we are excited to join in the community discussion at the conference.

VLDB kicks off this year with keynote talks from two speakers who provide a broad view of the landscape of data-driven technologies and their applications: Tim O'Reilly of O'Reilly Media, and David Campbell of Microsoft. We're delighted they were able to set aside time from their schedules to join us.

VLDB 2011 marks the end of the first year of a new experiment in paper reviewing and conference presentations designed by the VLDB Endowment. The Research Track of this year's conference consists of the full set of papers that appeared in the 2011 volume of the Proceedings of the VLDB Endowment (PVLDB). PVLDB is a journal consisting of short "conference-style" papers, with a monthly process designed to promote timely submission, review, and revision of scholarly results. We are grateful to the community of authors, reviewers and conference officers who worked year-round to make this experiment a reality. We also thank VLDB Proceedings Chair Uwe Roehm and PVLDB Editor-in-Chief H. V. Jagadish for their coordination and guidance in handling the rolling publication of PVLDB papers from the Research Track, as well as the final conference proceedings.

As in previous years, VLDB 2011 includes technical papers in research and industrial tracks, as well as a strong set of workshops, tutorials, demos and panels. In addition, as is traditional, selection committees singled out two research papers for recognition. The 10-year Best PaperAward for the paper from VLDB 2001 that has had the largest impact is "Generic Schema Matching with Cupid" by Jayant Madhavan, Philip A. Bernstein, and Erhard Rahm. The VDLB 2011 Best Paper Award is "RemusDB: Transparent High-Availability for Database Systems" by Umar Farooq Minha, Shriram Rajagopalan, Brendan Cully, Ashraf Aboulnaga, Ken Salem, and Andrew Warfield. We are grateful to the 10-year award selection committee consisting of Divesh Srivastava, Bettina Kemme, and Tova Milo; and the best paper award selection committee consisting of Umesh Dayal, Donald Kossman, and Stan Zdonik for their contribution.

This year's conference includes a Challenges and Vision Track held in coordination with the Computing Community Consortium (CCC). The track emphasizes visionary ideas, long term challenges, and opportunities in data-centric research that are outside of the current mainstream topics of the field. Submissions were judged on the extent to which they expand the possibilities and horizons of the field. The CCC sponsored three awards for best papers in the Challenges and Vision Track that will be announced at the conference.

This year's edition of VLDB received 553 submissions to the research track, of which 100 papers (18.1%) were accepted. We received 298 submissions (54%) spread from April 1, 2010 to February 1, 2011 for an average of 27 papers per month plus 255 submissions (46%) for the March 1, 2011 deadline. We would like to encourage the community to submit their work more evenly throughout the year.

The industrial and applications track received 41 submissions, of which 12 papers (27%) were accepted. The demonstrations track received 89 submissions, of which 28 (31 %) were accepted.

Our sincere thanks to all our colleagues who volunteered to serve as conference officers, program committee members and reviewers on this year's VLDB and PVLDB. We want to especially acknowledge the Research Track Co-Chairs -- Wolfgang Lehner, Sunita Sarawagi and Nick Koudas -- who took on the relentless task of driving the PVLDB review process to a conclusion every month. We thank Berthold Reinwald and Phil Bohannon for handling the industry and applications track; Jignesh Patel and Masatoshi Yoshikawa for handling the demonstrations program; and Qiong Luo and Gerome Miklau for handling the tutorials program.

We are also extremely grateful to General Chair Paul Larson and PVLDB Editor-in-Chief H.V. Jagadish, who calmly guided us all through this year's challenging new format. The staff at Microsoft's Conference Management Toolkit deserves recognition for their help in customizing CMT to accommodate the new structure of PVLDB and the VLDB conference.

Joseph M. Hellerstein José Blakeley

GENERAL INFORMATION

Name Badges

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

Note to Speakers

If you are scheduled to present, please ensure to check in with the Session Chair and test the projection setup well before the presentation time. Please upload your presentation to the in-room laptop to minimize switching time. We have conference volunteers on duty in each presentation room throughout the conference who can help you prepare your presentation files or alert the on-site technical support personnel if needed.

Mobile Phones, Pagers & Laptop Sound

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

Electricity

Electricity in United States of America / USA is 120 Volts, alternating at 60 cycles per second. If you travel to United States of America / USA with a device that does not accept 120 Volts at 60 Hertz, you will need a voltage converter.

Tipping

Tipping in restaurants is customary in the United States; a 15% tip is typical, with higher tips given for exceptional service. Tips are usually not included in the restaurant bill with the exception of large parties (typically six or more people). If the tip is included, the breakdown of the bill will read "gratuity".

Emergencies

Call 911 if you need urgent medical treatment or in case of emergency. You can call for an ambulance, police or firefighters using this number. You can dial 911 free of charge from any public phone.

VLDB WELCOME RECEPTION

Seattle Westin Hotel, Grand 3 Monday, August 29 6:00 – 9:00 PM

MICROSOFT RECEPTION

Seattle Westin Hotel, Grand 3 Tuesday, August 30 6:30 – 9:00 PM

VLDB BANQUET

Museum of Flight Wednesday, August 31 6:00 – 10:00 PM





One banquet ticket is included with your full or student VLDB conference registration. Additional banquet tickets for guests or "workshop only" participants may be purchased for \$100 per person. Tickets include: round-trip transportation via charter bus from Westin Hotel to Museum, a pre-dinner reception, and banquet dinner.

The VLDB Banquet will take place in the Great Gallery of The Museum of Flight. The museum captures the story of flight from the dawn of aviation to the Space Age and houses 54 of the world's most historic airplanes - authentic and in mint condition. Dozens of full-size aircraft are suspended from the ceiling of the steel and glass Great Gallery and appear to fly overhead in formation. At ground level visitors can examine up close such intriguing airplanes as the Blackbird spy plane and America's first presidential jet — the original Air Force One. The Museum of Flight is one of the largest air and space museums in the world, attracting more than 400,000 visitors annually.

Banquet Schedule

6:00 - 6:30 Chartered bus coaches depart every few minutes from the Westin Hotel.

Note: Buses will load on the Westlake Ave. side of the hotel, close to Valet Parking.

- 6:30 7:30 Reception and Viewing Exhibits at Museum of Flight
- 7:30 10:00 Banquet (return bus transportation provided to the Westin Hotel)







Keynote Talk 1: Towards a Global Brain Tim O'Reilly (O'Reilly Media)

Date: Tuesday, August 30, 2011 Time: 8:45 AM-10:00 AM Location: Grand 1 & 2

Tim O'Reilly (O'Reilly Media) Tim O'Reilly is the founder and CEO of O'Reilly Media Inc.,

thought by many to be the best computer book publisher in the world. O'Reilly Media also hosts conferences on technology topics, including the O'Reilly Open Source Convention, the Web 2.0 Summit, Strata: The Business of Data, and many others. O'Reilly's Make: magazine and Maker Faire has been compared to the West Coast Computer Faire, which launched the personal computer revolution. Tim's blog, the O'Reilly Radar "watches the alpha geeks" to determine emerging technology trends, and serves as a platform for advocacy about issues of importance to the technical community. Tim is also a partner at O'Reilly AlphaTech Ventures, O'Reilly's early stage venture firm, and is on the board of Safari Books Online.



Abstract. At the same time as we are seeing breakthrough after breakthrough in artificial intelligence, were also seeing the fulfillment of the vision of Vannevar Bush, JCR Licklider, and Doug Engelbart that computers could augment human information retrieval and problem solving. AI turned out not to be a matter of developing better algorithms, but of having enough data. The key applications of the web combine machine learning algorithms with techniques for harnessing the collective intelligence of users as captured in massive, interlinked cloud databases. Bit by bit, this is leading us towards a new kind of global brain, in which we have met the AI, and it is us. We and our devices are its senses, our databases are its memory, its habits, and even its dreams. This global brain is still a child, but as its parents, we have a responsibility to think about how best to raise it. What should we be teaching our future augmented selves? How can we make the emerging global consciousness not only more resilient, but more moral?

Keynote Talk 2: Is it still "Big Data" if it fits in my pocket? David Campbell (Microsoft)

Date: Wednesday, August 31, 2011 Time: 9:00 AM-10:00 AM

Location: Grand 1 & 2

David Campbell (Microsoft) David Campbell is a Microsoft Technical Fellow working in

Microsoft Corp.'s Server and Tools Business. Campbell joined Microsoft in 1994 from Digital Equipment Corp. as Microsoft began its push to become a credible enterprise software vendor. His early work at Microsoft included creating an OLE DB interface over the existing SQL Server storage engine, which helped to bootstrap SQL Server's present-generation query processor. He also worked closely with Mohsen Agsen, another Microsoft Technical Fellow, and the Microsoft Transaction Server team to add distributed transaction support to SQL Server 6.5. Microsoft made a bold move to re-architect SQL Server for the SQL Server 7.0 release. As a key technical member of the storage engine team, Campbell implemented the SQL Server lock manager



and other critical concurrency control mechanisms. He also implemented row-level locking in SQL Server 7.0, one of the hallmark features of the release.

Through the SQL Server 2000 and SQL Server 2005 releases, Campbell served in a variety of roles including product-level architect and general manager of product development. After the SQL Server 2005 release, he led a small team in redesigning SQL Server product development methodology. The new process, used to produce SQL Server 2008, resulted in SQL Server 2008 having the highest initial quality levels of any SQL Server release to date. As of August 2010, Campbell is serving as general manager of Microsoft's Data and Modeling Group, which oversees Microsoft's data modeling and data access strategies.

Campbell holds a number of patents in the data management, schema and software quality realms. He is also a frequent speaker at industry and research conferences on a variety of data management and software development topics. His current product development interests include cloud-scale computing, realizing value from ambient data, and multidimensional, context-rich computing experiences.

Abstract. "Big Data" is a hot topic but, in many ways, we are still trying to define what the phrase "Big Data" means. For many, there are more questions than answers at this point. Is it about size alone? complexity? variability? data shape? price/performance? new workloads? new types of users? Are existing data models, data management systems, data languages, and BI/ETL tools relevant in this space? Is MapReduce really a "major step backwards"? I have spent time over the last several years trying to answer many of these questions to my own satisfaction. As part of the journey I have witnessed a number of natural patterns that emerge in big data processing. In this talk I will present a catalog of these patterns and illustrate them across a scale spectrum from megabytes to 100s of petabytes. Finally, I will offer some thoughts around a systems and research agenda for this new world.

 10 Year Best Paper Award Keynote 1: Generic Schema Matching, Ten Years Later Generic Schema Matching With Cupid
 Philip A. Bernstein (Microsoft Research)

Jayant Madhavan (Google)

Erhard Rahm (University of Leipzig)

Date: Thursday, September 1, 2011 Time: 9:15 AM–10:00 AM Location: Grand 1 & 2

Abstract: In a paper published in the 2001 VLDB Conference, we proposed treating generic schema matching as an independent problem. We developed a taxonomy of existing techniques, a new schema matching algorithm, and an approach to comparative evaluation. Since then, the field has grown into a major research topic. We briefly summarize the new techniques that have been developed and applications of the techniques in the commercial world. We conclude by discussing future trends and recommendations for further work.

Philip A. Bernstein is a Principal Researcher at Microsoft Corporation. Over the past 35 years, he has been a product architect at Microsoft and Digital Equipment Corp., a professor at Harvard University and Wang Institute of Graduate Studies, and a VP Software at Sequoia Systems. During that time, he has published over 150 papers and two books on the theory and implementation of database systems, especially on transaction processing and metadata management. The second edition of his book "Transaction Processing" with Eric Newcomer was published in June 2009. His latest work focuses on database systems for cloud computing, on web search over structured data, and on object-to-relational mappings. He is an Editor-in-Chief of the VLDB



Journal, an ACM Fellow, a winner of the ACM SIGMOD Innovations Award, and a member of the Washington State Academy of Sciences and the National Academy of Engineering. He received a B.S. degree from Cornell and M.Sc. and Ph.D. from University of Toronto.

Jayant Madhavan is a Staff Software Engineer at Google Inc. where he is a member of the Structured Data Research group. His current research interests include exploring structure implicit in data on the Web, enabling new types of data collaboration and building interactive visualization over large datasets. He is the technical lead for the Google Fusion Tables product, a cloud data management solution, and had earlier led Google's deep-web crawler team. He was the Chief Architect at Transformic Inc., a data integration portal that was acquired by Google. He has served as a PC member for the AAAI, CIKM, EDBT, ICDE, IJCAI, SIGMOD, VLDB, and WWW conferences. He received a B.Tech. from IIT Bombay and M.S. and Ph.D. from the University of Washington. The VLDB 2001 paper was his first published research paper.



Erhard Rahm is a full professor for computer science at the University of Leipzig, Germany. He chairs the database group as well as the university's WDI lab focussing on research and development of web data integration solutions. His Ph.D. and habilitation degrees are from the University of Kaiserslautern, Germany. He has been a post-doc at IBM Research in Hawthorne, NY, and spent sabbaticals at Microsoft Research in Redmond, WA. He has organized several conferences and served as an editor of the VLDB Journal and on numerous program committees. He published about 150 peer-reviewed research papers and authored or co-edited more than ten books, including the 2011 Springer book on "Schema Matching and Mapping". His current research interests include data integration, metadata management, data management for the life sciences, and cloud data management.



| Mon | Monday August 29, 2011 | | | | | | | | | |
|-----------------|---------------------------|---|-----------------|---------------------------|------------|--|--|--|--|--|
| | Cascade 1BC | Cascade 2 | Cascade 1A | Vashon | St. Helens | | | | | |
| 08:00- 08:30 | | Breakfast-Cascade North & San Juan Foyers | | | | | | | | |
| | Cascade 1BC | Cascade 1BC Cascade 2 | | Vashon | St. Helens | | | | | |
| 08:30- 10:00 | ТРСТС | DBRank/MUD DBPL QDB (starts 9:00) | | QDB (starts 9:00) | DMSN | | | | | |
| 10:00- 10:30 | с | offee Break-Ca | scade North & S | San Juan Foyer | s | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Vashon | St. Helens | | | | | |
| 10:30- 12:00 | ТРСТС | DBRank/MUD | DBPL | QDB | DMSN | | | | | |
| 12:00- 01:30 | | Lunc | h Provided-Grai | nd 3 | | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Vashon | St. Helens | | | | | |
| 01:30- 03:00 | TPCTC DBRank/M | | DBPL | QDB | DMSN | | | | | |
| 03:00- 03:30 | С | offee Break-Ca | scade North & S | San Juan Foyer | s | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Vashon | St. Helens | | | | | |
| 03:00- 05:00 | TPCTC (ends 6:00) | DBRank/MUD (ends 5:30) | DBPL | QDB (ends 4:30) | DMSN | | | | | |
| 06:00- 09:00 | Welcome Reception-Grand 3 | | | | | | | | | |

| Tues | sday August 30, 2011 | | | | | | | | | |
|-----------------|---|---------------------------------------|--|--|----------------------------|------------------------------|---|--|--|--|
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | |
| 08:00- 08:30 | Breakfast-Grand Foyer | | | | | | | | | |
| 08:30- 08:45 | Opening Ceremony, Grand 1&2 | | | | | | | | | |
| 08:45- 10:00 | | | KEYNOT | E 1: Tim O'Reil | ly, Grand 1&2 | | | | | |
| 10:00- 10:30 | | | Cof | fee Break-Gra | nd Foyer | | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | |
| 10:30- 12:00 | Research 1: Distributed Systems | Research 2: Entity Matching | | Industrial 1: Testing, Debugging & Analysis | Research 3: Web | Tutorial 1: Bl | Demos A: Information Integration and Retrieval | | | |
| 12:00- 01:30 | | | Lu | nch Provided-(| Grand 3 | · | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | |
| 01:30- 03:30 | Research 4: GeoSpatial | Research 5: Uncertain Data | Research 6: Database Design | Research 7: Query Processing | Challenges and Vision 1 | Tutorial 2: Flash Storage | Demo B: Modern Hardware, Streaming and Benchmarking | | | |
| 03:30- 04:00 | | | Cot | fee Break-Gra | nd Foyer | | | | | |
| | Grand 1 Grand 2 Cascade 2 Fifth Avenue Vashon Cascade | | | | | Cascade 1BC | Grand Crescent | | | |
| 04:00- 06:00 | Research 8: Graph Data | Research 9: New HW Architecture | Research 10: Causality, Quality, Dependencies | Panel 1: Data Mgmt for Global Health | Challenges and Vision 2 | Tutorial 2: Flash Storage | Demos C: MapReduce, Crowdsourcing and Mining | | | |
| 06:30- 09:00 | Microsoft Reception-Grand 3 | | | | | | | | | |

| Wed | /ednesday August 31, 2011 | | | | | | | | | | |
|-----------------|--|--|-------------------------|---|---------------------------------------|------------------------------|---|--|--|--|--|
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | | |
| 08:00- 08:30 | Breakfast-Grand Foyer | | | | | | | | | | |
| 08:30- 09:00 | Reports: PC Chairs, VLDB 2012 General Chair, VLDB Endowment, Grand 1&2 | | | | | | | | | | |
| 09:00- 10:00 | | KEYNOTE 2: Dave Campbell, Grand 1&2 | | | | | | | | | |
| 10:00- 10:30 | | | Cot | fee Break-Gran | d Foyer | | | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | | |
| 10:30- 12:00 | Research 11: Graph Data | Research 12: Cloud Computing and High Availability | | Industrial 2: Large Scale- Analytics | Research 13: HCI | Tutorial 3: Deep Web | Demo: B Modern Hardware, Streaming and Benchmarking | | | | |
| 12:00- 01:30 | | | Lu | inch Provided-G | rand 3 | | | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | | |
| 01:30- 03:00 | Research 14: Integrity Maintenance | Research 15: Distributed Systems | | Industrial 3: Techniques for Large- Scale Data Management | Research 16: Streams and Events | Tutorial 4: Crowdsourcing | Demo C: MapReduce, Crowdsourcing and Mining | | | | |
| 03:00- 03:30 | | | Cot | fee Break-Gran | d Foyer | | | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescent | | | | |
| 03:30- 05:30 | Research 17: Privacy and Protection | Research 18: MapReduce and Hadoop | Research 19: Ranking | Research 20: Statistical Methods | Panel 2: Maximizing Impact | Tutorial 4: Crowdsourcing | Demo A: Information Integration and Retrieval | | | | |
| 06:00- 06:30 | | Buses depart | | | | | | | | | |
| 06:30- 07:30 | Reception and Viewing Time, Museum of Flight | | | | | | | | | | |
| 07:30- 10:30 | Banquet-Museum of Flight | | | | | | | | | | |

| Thu | sday September 1, 2011 | | | | | | | | |
|-----------------|--|---|----------------------|--|---|--|---|--|--|
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescer | | |
| 08:00- 08:30 | Breakfast-Grand Foyer | | | | | | | | |
| 08:30- 10:00 | VLDB AWARDS Including Invited Talk by 10-year Best Paper Award Winners Grand 1&2 | | | | | | | | |
| 10:00- 10:30 | | Coffee Break-Grand Foyer | | | | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescer | | |
| 10:30- 12:00 | Research 21: Graph Data | Research 22: Data Integration | PhD Workshop 1 | Research 23: Social Networks | | Tutorial 5: Graph Data Management | Demo C: MapReduce, Crowdsourcing and Mining | | |
| 12:00- 01:30 | | | L | unch Provided | -Grand 3 | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescer | | |
| 01:30- 03:00 | Research 24: Search and Ranking | Research 25: Statistical Methods | PhD Workshop 2 | Industrial 4: Large-Scale Distributed Systems | Research 26: Recommender Systems | Tutorial 6: Information in Social Nets | Demo B: Modern Hardware, Streaming and Benchmarking | | |
| 03:00- 03:30 | | | Со | ffee Break-Gra | and Foyer | | | | |
| | Grand 1 | Grand 2 | Cascade 2 | Fifth Avenue | Vashon | Cascade 1BC | Grand Crescer | | |
| 03:30- 05:30 | Research 27: GPU-based Architectures and Column- Store Indexing | Research 28: Transaction Processing | PhD Workshop 3 | Research 29: Web Data | Research 30: Skyline and String Matching | Tutorial 6: Information in Social Nets | Demo A: Information Integration and Retrieval | | |

VLDB 2011

| Frida | riday September 2, 2011 | | | | | | | | |
|-----------------|--|-------------------------------|----------------------|-------------------------------|--------|------------------------------|--|--|--|
| | Cascade 1BC | Cascade 2 | Cascade 1A | Fifth Avenue | Vashon | St. Helens | | | |
| 08:00- 08:30 | Breakfast-Cascade North & San Juan Foyers | | | | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Fifth Avenue | Vashon | St. Helens | | | |
| 08:30- 10:00 | PersDB | LADIS (starts 8:45) | SDM (starts 9:00) | BIRTE | VLDS | ADMS (starts 9:00) | | | |
| 10:00- 10:30 | Coffee Break-Cascade North & San Juan Foyers | | | | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Fifth Avenue | Vashon | St. Helens | | | |
| 10:30- 12:00 | PersDB | LADIS (10:15-12:15) | SDM | BIRTE (10:30-12:15) | VLDS | ADMS (10:45-12:15) | | | |
| 12:00- 01:30 | | | Lunch Provid | ded-Grand 3 | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Fifth Avenue | Vashon | St. Helens | | | |
| 01:30- 03:00 | PersDB | LADIS | SDM | BIRTE (1:30-3:15) | VLDS | ADMS (1:45-3:15) | | | |
| 03:00- 03:30 | Coffee Break-Cascade North & San Juan Foyers | | | | | | | | |
| | Cascade 1BC | Cascade 2 | Cascade 1A | Fifth Avenue | Vashon | St. Helens | | | |
| 03:30- 05:00 | PersDB (ends 5:30) | LADIS (ends 4:30) | SDM | BIRTE (3:45-5:00) | VLDS | ADMS (ends 5:30) | | | |

| Saturday Sept. 3, 2011 | | | | | | | | |
|------------------------|-----------------|-----------------|-----------------|-------------------|-----------------|-----------------|-------------|--|
| | Cascade 2 | | | | | | | |
| 08:00- 09:00 | 09:00- 10:00 | 10:00- 10:30 | 10:30- 12:00 | 12:00- 01:30 | 01:30- 03:00 | 03:00- 03:30 | 03:30-04:30 | |
| Breakfast | LADIS | Coffee Break | LADIS | Lunch Provided | LADIS | Coffee Break | LADIS | |



Tuesday 8:30 AM – 8:45 AM

Opening Ceremony

Location: Grand 1 & 2

Tuesday 8:45 AM - 10:00 AM

Keynote 1

Location: Grand 1 & 2 Towards a Global Brain

Tim O''Reilly (O'Reilly Media)

Tuesday 10:30 AM - 12:00 PM

Research Session 1: Distributed Systems

Location: Grand 1

Session Chair: Fei Xu

Distributed Threshold Querying of General Functions by a Difference of Monotonic Representation

Guy Sagy (Technion), Daniel Keren (Haifa University), Izchak Sharfman (Technion), Assaf Schuster (Technion)

Distributed Inference and Query Processing for RFID Tracking and Monitoring

Zhao Cao (University of Massachusetts), Charles Sutton (University of Edinburgh), Yanlei Diao (University of Massachusetts), Prashant Shenoy (University of Massachusetts)

Where in the World is My Data?

Sudarshan Kadambi (Bloomberg), Jianjun Chen (Yahoo!), Brian Cooper (Google), David Lomax (Yahoo!), Raghu Ramakrishnan (Yahoo!), Adam Silberstein (Yahoo!), Erwin Tam (Yahoo!), Hector Garcia-Molina (Stanford University)

Tuesday 10:30 AM – 12:00 PM

Research Session 2: Entity Matching

Location: Grand 2

Session Chair: Anish Das Sarma

Entity Matching: How Similar Is Similar

Jiannan Wang (Tsinghua University), Guoliang Li (Tsinghua University), Jeffrey Xu Yu (Chinese University of Hong Kong), Jianhua Feng (Tsinghua University)

Large-Scale Collective Entity Matching

Vibhor Rastogi (Yahoo! Research), Nilesh Dalvi (Yahoo!), Minos Garofalakis (Technical University of Crete)

Linking Temporal Records

Pei Li (University of Milan - Bicocca), Xin Dong (AT&T Labs), Andrea Maurino (University of Milan - Bicocca), Divesh Srivastava (AT&T Labs)

Tuesday 10:30 AM - 12:00 PM

Research Session 3: Web

Location: Vashon

Session Chair: Thomas Neumann

Output URL Bidding

Panagiotis Papadimitriou (Stanford University), Hector Garcia-Molina (Stanford University), Ali Dasdan (Ebay Inc), Santanu Kolay (Ebay Inc)

Automatic Wrappers for Large Scale Web Extraction

Nilesh Dalvi (Yahoo!), Ravi Kumar (Yahoo!), Mohamed Soliman (U. of Waterloo)

Recovering Semantics of Tables on the Web

Petros Venetis (Stanford University), Alon Halevy (Google), Jayant Madhavan (Google Inc), Marius Pasca (Google Inc), Warren Shen (Google), Fei Wu (Google Inc), Gengxin Miao (University of California, Santa Barbara), Chung Wu (Google Inc)

Tuesday 10:30 AM – 12:00 PM

Industrial Session 1: Database Systems Testing, Debugging, and Analysis

Location: Fifth Avenue

Session Chair: Yujun Wang

Consistent Synchronization Schemes for Workload Replay

Konstantinos Morfonios (Oracle), Romain Colle, Leonidas Galanis (Oracle), Supiti Buranawatanachoke (Oracle), Benoît Dageville (Oracle), Karl Dias (Oracle), Yujun Wang (Oracle)

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

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

HIWAS: Enabling Technology for Analysis of Clinical Data in XML Documents Joshua Hui (IBM Research - Almaden), Sarah Knoop (IBM Research - Almaden), Peter Schwarz (IBM Research - Almaden)

Tuesday 10:30 AM – 12:00 PM

Tutorial 1: New Frontiers in Business Intelligence Location: Cascade 1BC

New Frontiers in Business Intelligence

Surajit Chaudhuri and Vivek Narasayya

Tuesday 10:30 AM – 12:00 PM

Demo A: Information Integration and Information Retrieval Location: Grand Crescent

Demo presentations listed on pages 46-47

Tuesday 1:30 PM – 3:30 PM

Research Session 4: GeoSpatial

Location: Grand 1

Session Chair: Sang Kyun Cha

Graph Indexing of Road Networks for Shortest Path Queries with Label Restrictions Michael Rice (UCR), Vassilis Tsotras (UCR)

Efficient Processing of Top-k Spatial Preference Queries

João Rocha-Junior (NTNU), Akrivi Vlachou (NTNU), Christos Doulkeridis (NTNU), Kjetil Norvag (NTNU)

SXPath - Extending XPath towards Spatial Querying on Web Documents

Ermelinda Oro (DEIS-UNICAL, Altilia srl), Massimo Ruffolo (ICAR-CNR, Altilia srl), Steffen Staab (Institute WeST, University of Koblenz-Landau)

Efficient Algorithms for Finding Optimal Meeting Point on Road Networks

Da Yan (HKUST), Zhou Zhao (HKUST), Wilfred Ng (HKUST)

Tuesday 1:30 PM – 3:30 PM

Research Session 5: Uncertain Data

Location: Grand 2

Session Chair: Nilesh Dalvi

A Generic Framework for Handling Uncertain Data with Local Correlations Xiang Lian (HKUST), Lei Chen (HKUST)

Efficient Probabilistic Reverse Nearest Neighbor Query Processing on Uncertain Data

Thomas Bernecker (Ludwig-Maximilians-University), Tobias Emrich (Ludwig-Maximilians-University), Hans-Peter Kriegel (Ludwig-Maximilians-University), Matthias Renz (Ludwig-Maximilians-University), Stefan Zankl (Ludwig-Maximilians-University), Andreas Züfle (Ludwig-Maximilians-University)

Queries with Difference on Probabilistic Databases

Sanjeev Khanna (University of Pennsylvania), Sudeepa Roy (University of Pennsylvania), Val Tannen (University of Pennsylvania)

Optimizing Probabilistic Query Processing on Continuous Uncertain Data

Liping Peng (UMass Amherst), Yanlei Diao (University of Massachusetts), Anna Liu (UMass Amherst)

Tuesday 1:30 PM – 3:30 PM

Research Session 6: Database Design Location: Cascade 2

Session Chair: Wolfgang Lehner

CRIUS: User-Friendly Database Design

Li Qian (University of Michigan), Kristen LeFevre (University of Michigan), H. Jagadish (University of Michigan)

CoPhy: A Scalable, Portable, and Interactive Index Advisor for Large Workloads Debabrata Dash (ArcSight), Neoklis Polyzotis (UC Santa Cruz), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne (EPFL))

Compression Aware Physical Database Design

Hideaki Kimura (Brown University), Vivek Narasayya (Microsoft Research), Manoj Syamala (Microsoft Research)

Summary Graphs for Relational Database Schemas

Xiaoyan Yang (National University of Singapore), Cecilia Procopiuc (AT&T Labs), Divesh Srivastava (AT&T Labs)

Tuesday 1:30 PM – 3:30 PM

Research Session 7: Query Processing

Location: Fifth Avenue

Session Chair: Johann-Christoph Freytag

Generating Efficient Execution Plans for Vertically Partitioned XML Databases

Patrick Kling (University of Waterloo), M. Tamer Özsu (University of Waterloo), Khuzaima Daudjee (University of Waterloo)

Similarity Join Size Estimation using Locality Sensitive Hashing

Hongrae Lee (University of British Columbia), Raymond Ng (University of British Columbia), Kyuseok Shim (Seoul National University)

Accelerating Queries with Group-By and Join by Groupjoin

Guido Moerkotte (University of Mannheim), Thomas Neumann (Technische Universität München)

Optimizing Query Answering under Ontological Constraints

Giorgio Orsi (Oxford University), Andreas Pieris (Oxford University)

Tuesday 1:30 PM - 3:30 PM

Challenges and Vision 1

Location: Vashon

Session Chair: Hank Korth

Data Markets in the Cloud: An Opportunity for the Database Community

Magdalena Balazinska (University of Washington), Bill Howe (University of Washington), Dan Suciu (University of Washington)

Data is Dead... Without What-if Models

Peter Haas (IBM Research - Almaden), Paul Maglio (IBM Research - Almaden), Patricia Selinger (IBM Research - Almaden), Wang-Chiew Tan (IBM Research - Almaden & UCSC)

21

Data Generation for Application-Specific Benchmarking Y.C. Tay (National University of Singapore)

Reverse Data Management

Alexandra Meliou (University of Washington), Wolfgang Gatterbauer (University of Washington), Dan Suciu (University of Washington)

Anthropocentric Data Systems

Peter Triantafillou (University of Patras)

Tuesday 1:30 PM – 3:30 PM

Tutorial 2: Flash Storage (continues to 5:30 PM)

Location: Cascade 1BC

System Co-Design and Data Management for Flash Devices Philippe Bonnet, Luc Bouganim, Ioannis Koltsidas, and Stratis D. Viglas

Tuesday 1:30 PM – 3:30 PM

Demo B: Modern Hardware, Streaming, and Benchmarking

Location: Grand Crescent

Demo presentations listed on pages 46-47

Tuesday 4:00 PM - 6:00 PM

Research Session 8: Graph Data

Location: Grand 1

Session Chair: Thomas Willhalm

On Triangulation-based Dense Neighborhood Graph Discovery

Nan Wang (National University of Singapore), Jingbo Zhang (National University of Singapore), Kian-Lee Tan (National University of Singapore), Anthony Tung (National University of Singapore)

Human-Assisted Graph Search: It's Okay to Ask Questions

Aditya Parameswaran (Stanford University), Anish Das Sarma (Yahoo! Research), Hector Garcia-Molina (Stanford University), Neoklis Polyzotis (UC Santa Cruz), Jennifer Widom (Stanford University)

On Querying Historical Evolving Graph Sequences

Chenghui Ren (The University of Hong Kong), Eric Lo (HK Polytechnic University), Ben Kao (The University of Kong Kong), Xinjie Zhu (The University of Kong Kong), Reynold Cheng (University of Hong Kong)

Mining Top-K Large Structural Patterns in a Massive Network

Feida Zhu (Singapore Management University), Qiang Qu (Peking University), David Lo (Singapore Management University), Xifeng Yan (UCSB), Jiawei Han (UIUC), Philip Yu (UIC)

Tuesday 4:00 PM – 6:00 PM

Research Session 9: New Hardware Architecture

Location: Grand 2

Session Chair: Alfons Kemper

HYRISE - A Main Memory Hybrid Storage Engine

Martin Grund (Hasso-Plattner-Institute), Jens Krüger (Hasso-Plattner-Institute), Hasso Plattner (Hasso-Plattner Institute), Alexander Zeier (Hasso-Plattner Institute), Philippe Cudre-Mauroux (MIT), Samuel Madden (MIT)

Fast Set Intersection in Memory

Bolin Ding (UIUC), Arnd Christian König (Microsoft Research)

Efficiently Compiling Efficient Query Plans for Modern Hardware Thomas Neumann (Technische Universität München)

PALM: Parallel Architecture-Friendly Latch-Free Modifications to B+ Trees on Many-Core Processors

Jason Sewall (Intel Corporation), Jatin Chhugani (Intel Corporation), Changkyu Kim (Intel Corporation), Nadathur Satish (Intel Corporation), Pradeep Dubey (Intel Corporation)

VLDB 2011

Tuesday 4:00 PM - 6:00 PM

Research Session 10: Causality, Quality, and Dependencies Location: Cascade 2

Session Chair: Raghav Kaushik

The Complexity of Causality and Responsibility for Query Answers and non-Answers

Alexandra Meliou (University of Washington), Wolfgang Gatterbauer (University of Washington), Katherine Moore (University of Washington), Dan Suciu (University of Washington)

Guided Data Repair

Mohamed Yakout (Purdue University), Ahmed Elmagarmid (Qatar Computing Research Institute), Jennifer Neville (Purdue University), Mourad Ouzzani (Purdue University), Ihab Ilyas (University of Waterloo)

Stratification Criteria and Rewriting Techniques for Checking Chase Termination

Sergio Greco (Università della Calabria), Francesca Spezzano (Università della calabria), Irina Trubitsyna (DEIS, Università della Calabria)

Completeness of Queries over Incomplete Databases

Simon Razniewski (FU Bozen), Werner Nutt (FU Bozen)

Tuesday 4:00 PM - 6:00 PM

Panel 1: Data Management for Global Health Location: Fifth Ave

Moderator: Tapan Parikh

Data Management for Meeting Global Health Challenges

Tapan S. Parikh (UC Berkeley), Kuang Chen (UC Berkeley), Lucky Gunasekara (Global Viral Forecasting Initiative), Alon Halevy (Google), Andy Kanter (Columbia University), Rowena Luk (Dimagi Inc.), Peter Speyer (University of Washington)

Tuesday 4:00 PM - 6:00 PM

Challenges and Vision 2

Location: Vashon

Session Chair: Laura Haas

Resiliency-Aware Data Management

Matthias Boehm (TU Dresden), Wolfgang Lehner (TU Dresden), Christof Fetzer (TU Dresden)

Guided Interaction: Rethinking the Query-Result Paradigm

Arnab Nandi (University of Michigan), H.V. Jagadish (University of Michigan)

The Researcher's Guide to the Data Deluge: Querying a Scientific Database in Just a Few Seconds

Martin Kersten (CWI), Stratos Idreos (CWI), Stefan Manegold (CWI), Erietta Liarou (CWI)

Exploring the Coming Repositories of Reproducible Experiments: Challenges and Opportunities

Juliana Freire (University of Utah), Philippe Bonnet (University of Copenhagen, Denmark), Dennis Shasha (New York University)

Databases will Visualize Queries too

Wolfgang Gatterbauer (University of Washington)

Tuesday 4:00 PM - 6:00 PM

Tutorial 2 (cont): System Co-Design and Data Managment for Flash Devices

Location: Cascade 1BC

System Co-Design and Data Management for Flash Devices

Philippe Bonnet, Luc Bouganim, Ioannis Koltsidas, and Stratis D. Viglas

Tuesday 4:00 PM - 6:00 PM

Demo C: MapReduce, Crowdsourcing, and Mining

Location: Grand Crescent

Demo presentations listed on pages 46-47

Wednesday 8:30 AM - 9:00 AM

Reports

Location: Grand 1 & 2 PC Chairs, VLDB 2012 General Chair, VLDB Endowment

Wednesday 9:00 AM - 10:00 AM

Keynote 2

Location: Grand 1 & 2

Is it Still "Big Data" if it Fits in My Pocket? David Campbell (Microsoft)

Wednesday 10:30 AM - 12:00 PM

Research Session 11: Graph Data Location: Grand 1

Session Chair: Tamer Özsu

gStore: Answering SPARQL Queries via Subgraph Matching

Lei Zou (Peking University), Jinhui Mo (Peking University), Lei Chen (Hong Kong University of Science and Technology, China), M. Tamer Özsu (University of Waterloo), Zhao Dongyan (Peking University)

Efficient Subgraph Search over Large Uncertain Graphs

Ye Yuan (Northeastern University, China), Guoren Wang (Northeastern University, China), Haixun Wang (Microsoft Research Asia), Lei Chen (HKUST)

Scalable SPARQL Querying of Large RDF Graphs

Jiewen Huang (Yale University), Daniel Abadi (Yale University), Kun Ren (Northwestern Polytechnical University, China)

Wednesday 10:30 AM – 12:00 PM

Research Session 12: Cloud Computing and High-Availability Location: Grand 2

Session Chair: Alan Fekete

Albatross: Lightweight Elasticity in Shared Storage Databases for the Cloud using Live Data Migration

Sudipto Das (UC Santa Barbara), Shoji Nishimura (NEC Corporation), Divyakant Agrawal (UC Santa Barbara), Amr El Abbadi (UC Santa Barbara)

iCBS: Incremental Cost-based Scheduling under Piecewise Linear SLAs

Yun Chi (NEC Laboratories, America), Hyun Moon (NEC Labs America), Hakan Hacigumus (NEC Labs America)

RemusDB: Transparent High-Availability for Database Systems (Best Paper)

Umar Farooq Minhas (University of Waterloo), Shriram Rajagopalan (University of British Columbia), Brendan Cully (University of British Columbia), Ashraf Aboulnaga (University of Waterloo), Ken Salem (University of Waterloo), Andrew Warfield (University of British Columbia)

Wednesday 10:30 AM – 12:00 PM

Research Session 13: Human-Computer Interaction

Location: Vashon

Session Chair: H.V. Jagadish

SnipSuggest: Context-Aware Autocompletion for SQL

Nodira Khoussainova (University of Washington), YongChul Kwon (University of Washington), Magdalena Balazinska (University of Washington), Dan Suciu (University of Washington)

A Probabilistic Approach for Automatically Filling Form-Based Web Interfaces

Guilherme Toda (Federal University of Amazonas), Eli Cortez (Federal University of Amazonas), Altigran da Silva (Federal University of Amazonas), Edleno de Moura (Federal University of Amazonas)

Query Expansion Based on Clustered Results

Ziyang Liu (Arizona State University), Sivaramakrishnan Natarajan (Arizona State University), Yi Chen (ASU)

Wednesday 10:30 AM - 12:00 PM

Industrial Session 2: Large-Scale Analytics

Location: Fifth Avenue

Session Chair: Nico Bruno

Bridging Two Worlds with RICE

Philipp Große (SAP), Wolfgang Lehner (TU Dresden), Thomas Weichert (SAP), Franz Färber (SAP), Wen-Syan Li (SAP)

Jaql: A Scripting Language for Large Scale Semistructured Data Analysis

Kevin Beyer (IBM), Vuk Ercegovac (IBM), Rainer Gemulla (Max-Planck-Institut für Informatik), Andrey Balmin (IBM Research - Almaden), Mohamed Eltabakh (IBM), Carl-Christian Kanne (IBM Research - Almaden), Fatma Ozcan (IBM Research - Almaden), Eugene Shekita (IBM Research - Almaden)

Evaluation Strategies for Top-k Queries over Memory-Resident Inverted Indexes Marcus Fontoura (Google Inc), Vanja Josifovski (Yahoo! Research), Jinhui Liu (Yahoo! Research), Srihari Venkatesan (Yahoo! Research), Xiangfei Zhu (Yahoo! Research), Jason Zien (Yahoo! Research)

Wednesday 10:30 AM - 12:00 PM

Tutorial 3: Exploration of Deep Web Repositories

Location: Cascade 1BC

Exploration of Deep Web Repositories Nan Zhang and Gautam Das

Wednesday 10:30 AM - 12:00 PM

Demo B: Modern Hardware, Streaming, and Benchmarking Location: Grand Crescent

Demo presentations listed on pages 46-47

Wednesday 1:30 PM - 3:00 PM

Research Session 14: Integrity Maintenance Location: Grand 1

Session Chair: Divesh Srivastava

Update Rewriting and Integrity Constraint Maintenance in a Schema Evolution Support System: PRISM++

Carlo Curino (MIT), Hyun Moon (NEC Labs America), Alin Deutsch (UCSD), Carlo Zaniolo (UCLA)

Business Policy Modeling and Enforcement in Databases

Ahmed Ataullah (University of Waterloo), Frank Tompa (University of Waterloo)

Data Coordination: Supporting Contingent Updates

Michael Lawrence (University of British Columbia), Rachel Pottinger (University of British Columbia), Sheryl Staub-French (University of British Columbia)

Wednesday 1:30 PM - 3:00 PM

Research Session 15: Distributed Systems

Location: Grand 2

Session Chair: Theo Härder

Using Paxos to Build a Scalable, Consistent, and Highly Available Datastore Jun Rao (LinkedIn), Eugene Shekita (IBM Research), Sandeep Tata (IBM Research)

A Framework for Supporting DBMS-like Indexes in the Cloud

Gang Chen (Zhejiang University), Hoang Tam Vo (School of Computing), Sai Wu (National Univ. of Singapore), Beng Chin Ooi (National University of Singapo), M. Tamer Özsu (University of Waterloo)

Serializable Snapshot Isolation for Replicated Databases in High-Update Scenarios Hyungsoo Jung (University of Sydney), Hyuck Han (Seoul National University), Alan Fekete (University of Sydney), Uwe Roehm (University of Sydney)

Wednesday 1:30 PM – 3:00 PM

Research Session 16: Streams and Events Location: Vashon

Session Chair: Themis Palpanas

Lightweight Graphical Models for Selectivity Estimation Without Independence Assumptions

Kostas Tzoumas (Aalborg University, Denmark), Amol Deshpande (University of Maryland), Christian Jensen (Aarhus University)

Active Complex Event Processing over Event Streams

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

Massive Scale-out of Expensive Continuous Queries

Erik Zeitler (Uppsala University), Tore Risch (Uppsala University)

Wednesday 1:30 PM – 3:00 PM

Industrial Session 3: Techniques for Large-Scale Data

Management

Location: Fifth Avenue

Session Chair: Sameh Elnikety

Online Expansion of Large-scale Data Warehouses

Jeffrey Cohen (EMC), John Eshleman (EMC), Brian Hagenbuch (EMC), Joy Kent (EMC), Christopher Pedrotti (EMC), Gavin Sherry (EMC), Florian Waas (EMC)

Auto-Grouping Emails For Faster E-Discovery

Sachindra Joshi (IBM Research India), Danish Contractor (IBM Research India), Kenney Ng (IBM Software Group, USA), Prasad Deshpande (IBM Research - India), Thomas Hampp (IBM Software Group, Germany)

Web Scale Taxonomy Cleansing

Taesung Lee (POSTECH), Zhongyuan Wang (Microsoft Research Asia), Haixun Wang (Microsoft Research Asia), Seung-won Hwang (POSTECH)

Wednesday <u>1:30 PM – 3:00 PM</u>

Tutorial 4: Crowdsourcing Applications and Platforms: A Data

Management Perspective (continues to 5:30 PM)

Location: Cascade 1BC

Crowdsourcing Applications and Platforms: A Data Management Perspective Anhai Doan, Michael Franklin, Donald Kossmann, and Tim Kraska

Wednesday 1:30 PM – 3:00 PM

Demo C: MapReduce, Crowdsourcing, and Mining

Location: Grand Crescent

Demo presentations listed on pages 46-47

Wednesday 3:30 PM – 5:30 PM

Research Session 17: Privacy and Protection

Location: Grand 1

Session Chair: Graham Cormode

Personalized Privacy Protection in Social Networks Mingxuan Yuan (HKUST), Lei Chen (HKUST), Philip Yu (UIC)

Publishing Set-Valued Data via Differential Privacy

Rui Chen (Concordia University), Noman Mohammed (Concordia University), Benjamin C. M. Fung (Concordia University), Bipin c. Desai (Concordia University), Li Xiong (Emory University)

Private Analysis of Graph Structure

Vishesh Karwa (Pennsylvania State University), Sofya Raskhodnikova (Pennsylvania State University), Adam Smith (Pennsylvania State University), Grigory Yaroslavtsev (Pennsylvania State University)

Surrogate Parenthood: Protected and Informative Graphs

Barbara Blaustein (MITRE), Adriane Chapman (MITRE), Len Seligman (MITRE), M. David Allen (MITRE), Arnon Rosenthal (MITRE)

Wednesday 3:30 PM – <u>5:30 PM</u>

Research Session 18: MapReduce and Hadoop Location: Grand 2

Session Chair: Jingren Zhou

Column-Oriented Storage Techniques for MapReduce

Avrilia Floratou (University of Wisconsin-Madison), Jignesh Patel (University of Wisconsin-Madison), Eugene Shekita (IBM Research), Sandeep Tata (IBM Research)

Automatic Optimization for MapReduce Programs

Eaman Jahani (University of Michigan), Michael Cafarella (University of Michigan), Christopher Ré (University of Wisconsin-Madison)

CoHadoop: Flexible Data Placement and Its Exploitation in Hadoop

Mohamed Eltabakh (IBM), Yuanyuan Tian (IBM Research - Almaden), Fatma Ozcan (IBM Research - Almaden), Rainer Gemulla (Max-Planck-Institut für Informatik), Aljoscha Krettek (IBM Germany), John McPherson (IBM Research - Almaden)

Profiling, What-if Analysis, and Cost-based Optimization of MapReduce Programs Herodotos Herodotou (Duke University), Shivnath Babu (Duke University)

Wednesday 3:30 PM – 5:30 PM

Research Session 19: Ranking

Location: Cascade 2

Session Chair: Ihab Ilyas

On Pruning for Top-K Ranking in Uncertain Databases

Chonghai Wang (University of Alberta), Li Yan Yuan (University of Alberta), Jia-Huai You (University of Alberta), Osmar Zaiane (University of Alberta), Jian Pei (Simon Fraser University)

PathSim: Meta Path-Based Top-K Similarity Search in Heterogeneous Information Networks

Yizhou Sun (UIUC), Jiawei Han (UIUC), Xifeng Yan (UCSB), Philip Yu (UIC), Tianyi Wu (Microsoft)

Optimizing and Parallelizing Ranked Enumeration

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

Efficient Rank Join with Aggregation Constraints

Min Xie (University of British Columbia), Laks Lakshmanan (University of British Columbia), Peter Wood (Birkbeck, University of London)

Wednesday 3:30 PM – 5:30 PM

Research Session 20: Statistical Methods

Location: Fifth Avenue

Session Chair: Rainer Gemulla

Tuffy: Scaling up Statistical Inference in Markov Logic Networks using an RDBMS

Feng Niu (University of Wisconsin), Christopher Ré (University of Wisconsin-Madison), AnHai Doan (University of Wisconsin), Jude Shavlik (University of Wisconsin-Madison)

Dissemination of Models over Time-Varying Data

Yongluan Zhou (University of Southern Denmark), Zografoula Vagena (Rice University), Jonas Haustad (University of Southern Denmark)

Storing Matrices on Disk: Theory and Practice Revisited

Yi Zhang (Duke University), Kamesh Munagala (Duke University), Jun Yang (Duke University)

Online Aggregation for Large MapReduce Jobs

Niketan Pansare (Rice University), Vinayak Borkar (UC Irvine), Chris Jermaine (University of Florida), Tyson Condie (Yahoo! Research)

Wednesday 3:30 PM – 5:30 PM

Panel 2: Maximizing Impact

Location: Vashon

Panel Discussion: Maximizing Impact

David DeWitt (Microsoft), Juliana Freire (NYU Polytechnic), Ed Lazowska (University of Washington), Sam Madden (MIT), Jennifer Widom (Stanford)

Wednesday 3:30 PM – 5:30 PM

Tutorial 4 (cont): Crowdsourcing Applications and Platforms: A Data Management Perspective

Location: Cascade 1BC

Crowdsourcing Applications and Platforms: A Data Management Perspective Anhai Doan, Michael Franklin, Donald Kossmann, and Tim Kraska

Wednesday 3:30 PM - 5:30 PM

Demo A: Information Integration and Information Retrieval Location: Grand Crescent

Demo presentations listed on pages 46-47

Thursday 8:30 AM – 10:00 AM

VLDB Awards & Invited Talk by 10-Year Best Paper Award Winners Location: Grand 1 & 2

10 Year Best Paper Award Keynote 1: Generic Schema Matching, Ten Years Later Philip A. Bernstein (Microsoft Research)

Jayant Madhavan (Google)

Erhard Rahm (University of Leipzig)

Thursday 10:30 AM - 12:00 PM

Research Session 21: Graph Data

Location: Grand 1

Session Chair: Michael Rys

Distance-Constraint Reachability Computation in Uncertain Graphs

Ruoming Jin (Kent State University), Lin Liu (Kent State University), Bolin Ding (UIUC), Haixun Wang (Microsoft Research Asia)

Keyword Search in Graphs: Finding r-cliques

Mehdi Kargar (York University), Aijun An (York University)

On Link-based Similarity Join

Liwen Sun (University of Hong Kong), Reynold Cheng (University of Hong Kong), Xiang Li (University of Hong Kong), David Cheung (University of Hong Kong), Jiawei Han (UIUC)

Thursday 10:30 AM – 12:00 PM

Research Session 22: Data Integration

Location: Grand 2

Session Chair: Rachel Pottinger

Synthesizing Products for Online Catalogs

Hoa Nguyen (University of Utah), Ariel Fuxman (Microsoft Research), Stelios Paparizos (Microsoft Research), Juliana Freire (University of Utah), Rakesh Agrawal (Microsoft Research)

Online Data Fusion

Xuan Liu (National Univ. of Singapore), Xin Dong (AT&T Labs), Beng Chin Ooi (National University of Singapo), Divesh Srivastava (AT&T Labs)

Thursday 10:30 AM - 12:00 PM

Research Session 23: Social Networks

Location: Fifth Avenue

Session Chair: Cong Yu

Structural Trend Analysis For Online Social Networks

Ceren Budak (UCSB), Divyakant Agrawal (University of California, Santa Barbara), Amr El Abbadi (UC Santa Barbara)

On Social-Temporal Group Query with Acquaintance Constraint

De-Nian Yang (Academia Sinica), Yi-Ling Chen (National Taiwan University), Wang-Chien Lee (The Penn State University), Ming-Syan Chen (National Taiwan University)

Social Content Matching in MapReduce

Gianmarco De Francisci Morales (IMT Lucca), Aristides Gionis (Yahoo! Research), Mauro Sozio (MPI)

Thursday 10:30 AM - 12:00 PM

PhD Workshop 1

Location: Cascade 2

Query Processing in a Self-Organized Storage System Hannes Mühleisen (Freie Universität Berlin)

Efficient Top-k Searching According to User Preferences Based on Fuzzy Functions

With Usage of Tree-Oriented Data Structures Matus Ondreicka (Charles University in Prague)

Top-k Web Service Compositions in the Context of User Preferences Karim Benouaret (University of Lyon)

Thursday 10:30 AM – 12:00 PM

Tutorial 5: Graph Data Management Systems for New Application Domains

Location: Cascade 1BC

Graph Data Management Systems for New Application Domains

Philippe Cudré-Mauroux (University of Fribourg), Sameh Elnikety (Microsoft Research)

Thursday 10:30 AM – 12:00 PM

Demo C: MapReduce, Crowdsourcing, and Mining

Location: Grand Crescent

Demo presentations listed on pages 46-47

Thursday 1:30 PM – 3:00 PM

Research Session 24: Searching and Ranking

Location: Grand 1

Session Chair: Frank Tompa

Fast Incremental and Personalized PageRank

Bahman Bahmani (Stanford University), Abdur Chowdhury (Twitter Inc.), Ashish Goel (Stanford University, Twitter Inc.)

Efficient Diversification of Web Search Results

Gabriele Capannini (ISTI CNR), Franco Maria Nardini (ISTI-CNR), Raffaele Perego (ISTI-CNR), Fabrizio Silvestri (ISTI-CNR)

Keyword Search on Form Results

Aditya Ramesh (Stanford University), S. Sudarshan (IIT Bombay), Purva Joshi (IIT Bombay)

Thursday 1:30 PM – 3:00 PM

Research Session 25: Statistical Methods

Location: Grand 2

Session Chair: Cesar Galindo-Legaria

Incrementally Maintaining Classification using an RDBMS

Mehmet Levent Koc (University of Wisconsin-Madison), Christopher Ré (University of Wisconsin-Madison)

An Incremental Hausdorff Distance Calculation Algorithm

Sarana Nutanong (University of Maryland), Edwin Jacox (University of Maryland), Hanan Samet (University of Maryland)

Structure-Aware Sampling: Flexible and Accurate Summarization

Edith Cohen (AT&T Labs), Graham Cormode (AT&T Labs), Nick Duffield (AT&T Labs)

Thursday 1:30 PM – 3:00 PM

Research Session 26: Recommender Systems

Location: Vashon

Session Chair: Wolfgang Gatterbauer

Personalized Social Recommendations - Accurate or Private?

Ashwin Machanavajjhala (Yahoo! Research), Aleksandra Korolova (Stanford University), Atish Das Sarma (Google)

RecBench: Benchmarks for Evaluating Performance of Recommender System Architectures

Justin Levandoski (University of Minnesota), Michael Ekstrand (University of Minnesota), Michael Ludwig (University of Minnesota), Ahmed Eldawy (University of Minnesota), Mohamed Mokbel (University of Minnesota), John Riedl (University of Minnesota)

MRI: Meaningful Interpretations of Collaborative Ratings

Mahashweta Das (University of Texas, Arlington), Sihem Amer-Yahia (Yahoo Research, USA), Gautam Das (University of Texas, Arlington), Cong Yu (Google Research)

Thursday 1:30 PM – 3:00 PM

PhD Workshop 2

Location: Cascade 2

Mixed Workload Management for In-Memory Databases Johannes Wust (Hasso Plattner Institute)

Scaling Web Applications: A Temporal Approach Zhiwu Xie (University of New Mexico)

Research on a Schema and Data Versioning System

Bob Wall (Montana State University)

Thursday 1:30 PM – 3:00 PM

Industrial Session 4: Large-Scale Distributed Systems Location: Fifth Avenue

Session Chair: Praveen Seshadri

Tenzing - A SQL Implementation on the MapReduce Framework

Biswapesh Chattopadhyay (Google), Liang Lin (Google), Weiran Liu (Google), Sagar Mittal (Google), Prathyusha Aragonda (Google), Vera Lychagina (Google), Younghee Kwon (Google), Michael Wong (Google)

33

An Algebraic Approach for Data-Centric Scientific Workflows Eduardo Ogasawara (COPPE/UFRJ), Jonas Dias (COPPE/UFRJ), Daniel de Oliveira (COPPE/

UFRJ), Fábio Porto (LNCC), Patrick Valduriez (INRIA), Marta Mattoso (COPPE/UFRJ)

Citrusleaf: A Real-Time NoSQL DB which Preserves ACID V. Srinivasan (Citrusleaf), Brian Bulkowski (Citrusleaf)

Thursday 1:30 PM - 3:00 PM

Tutorial 6: Information Diffusion In Social Networks: Observing and Influencing Societal Interests (continues to 5:30 PM)

Location: Cascade 1BC

Information Diffusion In Social Networks: Observing and Influencing Societal Interests

Divyakant Agrawal, Ceren Budak, and Amr El Abbadi

Thursday 1:30 PM – 3:00 PM

Demo B: Modern Hardware, Streaming, and Benchmarking Location: Grand Crescent

Demo presentations listed on pages 46-47

Thursday <u>3:30 PM – 5:30 PM</u>

Research Session 27: GPU-based Architectures and Column-Store Indexing

Location: Grand 1

Session Chair: Kenneth Ross

Fast Sparse Matrix-Vector Multiplication on GPUs: Implications for Graph Mining Xintian Yang (Ohio State University), Srinivasan Parthasarathy (Ohio State University),

Ponnuswamy Sadayappan (Ohio State University)

High-Throughput Transaction Executions on Graphics Processors

Bingsheng He (Nanyang Technological University), Jeffrey Xu Yu (Chinese University of Hong Kong)

Efficient Parallel Lists Intersection and Index Compression Algorithms using Graphics Processing Units

Naiyong Ao (Nankai University), Fan Zhang (Nankai University), Di Wu (Nankai University), Douglas Stones (Monash University), Gang Wang (Nankai University), Xiaoguang Liu (Nankai University), Jing Liu (Nankai University), Sheng Lin (Nankai University)

Merging What's Cracked, Cracking What's Merged: Adaptive Indexing in Main-Memory Column-Stores

Stratos Idreos (CWI), Stefan Manegold (CWI), Harumi Kuno (HP Labs), Goetz Graefe (HP Labs)

Thursday 3:30 PM – 5:30 PM

Research Session 28: Transaction Processing

Location: Grand 2

Session Chair: S. Sudarshan

PLP: Page Latch-free Shared-everything OLTP

Ippokratis Pandis (Carnegie Mellon University), Pinar Tozun (EPFL), Ryan Johnson (University of Toronto), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne (EPFL))

Implementing Performance Competitive Logical Recovery

David Lomet (Microsoft Research, USA), Kostas Tzoumas (Aalborg University, Denmark), Michael Zwilling (Microsoft)

Entangled Transactions

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

Optimistic Concurrency Control by Melding Trees

Philip Bernstein (Microsoft Research), Colin Reid (Microsoft Research), Ming Wu (Microsoft), Xinhao Yuan (Tsinghua University)

Thursday 3:30 PM – 5:30 PM

Research Session 29: Web Data

Location: Fifth Avenue

Session Chair: Xin Dong

Hyper-Local, Directions-Based Ranking of Places

Petros Venetis (Stanford University), Hector Gonzalez (Google Inc), Christian Jensen (Aarhus University), Alon Halevy (Google)

Optimal Schemes for Robust Web Extraction

Aditya Parameswaran (Stanford University), Nilesh Dalvi (Yahoo!), Hector Garcia-Molina (Stanford University), Rajeev Rastogi (Yahoo! India)

OXPath: A Language for Scalable, Memory-efficient Data Extraction from Web Applications

Tim Furche (Oxford University), Georg Gottlob (Oxford University), Giovanni Grasso (Oxford University), Christian Schallhart (Oxford University), Andrew Sellers (Oxford University)

Randomized Generalization for Aggregate Suppression over Hidden Web Databases

Xin Jin (George Washington U), Nan Zhang (George Washington U), Aditya Mone (UT Arlington), Gautam Das (UT Arlington)

Thursday 3:30 PM - 5:30 PM

Research Session 30: Skyline and String Matching Location: Vashon

Session Chair: Mohamed Mokbel

ZINC: Efficient Indexing for Skyline Computation

Bin Liu (National Univ of Singapore), Chee-Yong Chan (National University of Singapore)

QSkycube: Efficient Skycube Computation Using Point-Based Space Partitioning Jongwuk Lee (POSTECH), Seung-won Hwang (POSTECH)

A Subsequence Matching with Gaps-Range-Tolerances Framework: A Query-By-Humming Application

Alexios Kotsifakos (University of Athens), Panagiotis Papapetrou (Aalto University), Jaakko Hollmen (Aalto University), Dimitris Gunopulos (University of Athens)

35

Approximate Substring Matching over Uncertain Strings

Tingjian Ge (University of Kentucky), Zheng Li (University of Kentucky)

VLDB 2011

Thursday 3:30 PM – 5:30 PM

PhD Workshop 3

Location: Cascade 2

Matching Tree Patterns on Partial-trees Shachar Harussi (Tel Aviv University)

Knowledge-Based Complex Event Processing

Kia Teymourian (Free University Berlin)

Thursday 3:30 PM – 5:30 PM

Tutorial 6 (cont): Information Diffusion In Social Networks: Observing and Influencing Societal Interests

Location: Cascade 1BC

Information Diffusion In Social Networks: Observing and Influencing Societal Interests

Divyakant Agrawal, Ceren Budak, and Amr El Abbadi

Thursday 3:30 PM - 5:30 PM

Demo A: Information Integration and Information Retrieval Location: Grand Crescent

Demo presentations listed on pages 46-47



NOTES

Tutorial 1: New Frontiers in Business Intelligence

Surajit Chaudhuri and Vivek Narasayya

Date: Tuesday, August 30, 2011 Time: 10:30 AM-12:00 PM Room: Cascade 1BC

Abstract: Business intelligence (BI) software is a collection of decision support technologies for the enterprise aimed at enabling knowledge workers such as executives, managers and analysts to make better and faster decisions. When compared to the BI landscape of the mid-90s we observe that today's BI technology has progressed well beyond OLAP servers, parallel DBMSs, and classical ETL of that era. Several new frontiers in BI have emerged including "Big Data" engines, near real-time BI, predictive analytics, enterprise search and cloud data services. In this tutorial, we will first provide a broad overview of the current BI landscape identifying important use cases, describing key technologies, and highlighting relationships across these technologies. Next, we drill-down into the newly emerging frontiers of BI discussed above. For each area we discuss the state-of-the-art, highlight architectural considerations and present open research problems.

Surajit Chaudhuri is a Research Manager at Microsoft Research, Redmond. He started the AutoAdmin project on self-tuning database systems at Microsoft Research. Surajit has also worked in the area of data cleaning. His research on both physical database design and data cleaning has been incorporated in Microsoft products and services such as Microsoft SQL Server and Bing. Surajit received his Ph.D. from Stanford University and is an ACM Fellow. He was awarded the ACM SIGMOD Contributions award in 2004, a VLDB 10-year Best paper Award in 2007 and the ACM SIGMOD Innovations Award in 2011.



Vivek Narasayya is a Principal Researcher at Microsoft Research, Redmond. He is interested broadly in data management, focusing on the areas of self-tuning database systems, query processing, query optimization, and resource management in databases. He received his Ph.D. from the University of Washington, Seattle. He was awarded a VLDB 10-year Best paper Award in 2007.



Tutorial 2: System Co-Design and Data Management for Flash Devices

Philippe Bonnet, Luc Bouganim, Ioannis Koltsidas, and Stratis D. Viglas

Date: Tuesday, August 30, 2011 Time: 1:30 PM-6:00 PM Room: Cascade 1BC

Abstract: Flash devices are emerging as a replacement for disks. How does this evolution impact the design of data management systems? While flash devices have been available for years, this question is still open. In this tutorial, we share two views on the development of data management systems for flash devices. The first view considers that flash devices introduce so much complexity that it is necessary to reconsider the strictly layered approach between storage system, operating system and data management system. The second view considers that data management systems should recognize the complexity of flash devices and leverage the characteristics of different classes of devices for different usage patterns. Throughout the tutorial, we will cover the data management stack: from the fundamentals of flash technology, through storage for database systems and the manipulation of flash-resident data, to query processing.

Philippe Bonnet is associate professor at IT University of Copenhagen; his research interests lie in the area of data management and system performance. Philippe currently serves as chair for the SIGMOD repeatability committee. Philippe is co-author of a book on database tuning together with Dennis Shasha, NYU.



Luc Bouganim is research director at INRIA Paris-Rocquencourt and vicehead of the SMIS project, which focuses on secured and mobile information systems. Luc co-authored more than 50 international conference and journal papers. Since 2000, his research interests lie in the area of data management on specific hardware architecture, and more precisely on secure chips. Luc got best paper awards at VLDB in 2000 for his work on PicoDBMS, a database system for smartcard.

Ioannis Koltsidas is a Research Staff Member in the Storage Technologies Department of IBM Research – Zurich. He received a Ph.D. in Computer Science from the University of Edinburgh, UK, in 2010, and a B.Sc. degree in Electrical and Computer Engineering from the National Technical University of Athens, Greece, in 2006. His research interests lie in the areas of systems storage and database storage, with an emphasis on buffer management, caching and automatic tiering. His current focus is on high-performance, high-scalability storage systems employing solid-state storage technologies.

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

His research interests lie in the areas of query processing and optimization, data storage and indexing, data stream management, and distributed computing.



Tutorial 3: Exploration of Deep Web Repositories

Nan Zhang and Gautam Das

Date: Wednesday, August 31, 2011 Time: 10:30 AM-12:00 PM Room: Cascade 1BC

Abstract: With the proliferation of online repositories (e.g., databases or document corpora) hidden behind proprietary web interfaces, e.g., keyword-/form-based search and hierarchical/ graph-based browsing interfaces, efficient ways of exploring contents in such hidden repositories are of increasing importance.

There are two key challenges: one on the proper understanding of interfaces, and the other on the efficient exploration, e.g., crawling, sampling and analytical processing, of very large repositories. In this tutorial, we focus on the fundamental developments in the field, including web interface understanding, crawling, sampling, and data analytics over web repositories with various types of interfaces and containing structured or unstructured data. Our goal is to encourage audience to initiate their own research in these exciting areas.

Dr. Nan Zhang is an Assistant Professor of Computer Science at the George Washington University, Washington, DC, USA. Prior to ioining GWU, he was an assistant professor of Computer Science and Engineering at the University of Texas at Arlington from 2006 to 2008. He received the B.S. degree from Peking University in 2001 and the Ph.D. degree from Texas A&M University in 2006, both in computer science. His current research interests include databases and information security/privacy. He received the NSF CAREER award in 2008.



Gautam Das is a Full Professor in the Computer Science and Engineering Department of the University of Texas at Arlington. Prior to UTA, Dr Das has held positions at Microsoft Research, Compag Corporation and the University of Memphis, as well as visiting positions at IBM Research. He graduated with a BTech in computer science from IIT Kanpur, India, and with a PhD in computer science from the University of Wisconsin-Madison. Dr. Das's research interests span data mining, information retrieval, databases, approximate query processing, applied graph and network algorithms, and computational geometry. His research has resulted in over 100 papers, many of which have appeared in premier conferences and journals such as SIGMOD, VLDB,



ICDE, KDD, TODS, and TKDE, including several best paper awards. Dr. Das has served as the General Chair of ICIT 2009, Program Committee Vice-Chair of ICDM 2011, Program Chair of COMAD 2008, ICDE DBRank 2007, Best Paper Awards Chair of KDD 2006, Best Papers Awards committee of DAFSAA 2008, Program Chair of ICIT 2004, as well as in program committees of premier conferences such as SIGMOD, PODS, WWW, ICDE, KDD, and ICML, His research has been supported by grants from federal and state agencies such as the National Science Foundation, Office of Naval Research, Department of Education, Texas Higher Education Coordination Board, as well as industry such as Nokia, Microsoft, Cadence, and Apollo.

Tutorial 4: Crowdsourcing Applications and Platforms: A Data Management Perspective

Anhai Doan, Michael Franklin, Donald Kossmann, and Tim Kraska

Date: Wednesday, August 31, 2011 Time: 1:30 PM-6:00 PM Room: Cascade 1BC

Abstract:: Over the past decade, crowdsourcing has emerged as a major problem-solving and data-gathering paradigm on the World-Wide Web, and has attracted much interest in the database community. This tutorial seeks to spark further interest and contribute to the gathering momentum of crowdsourcing in the community. It provides an overview of current crowdsourcing platforms and applications, describes data management approaches to crowdsourcing, and identifies data management-specific research challenges in the area.

AnHai Doanis an Associate Professor of Computer Science at the University of Wisconsin, Madison, He has worked extensively in crowdsourcing, starting with the MOBS project in 2002-2006, and the Cimple/DBLife project from 2006 to date. In these projects he has studied how to crowdsource data integration problems, and the problems of building structured databases from unstructured Web data. Currently he is on leave, working as Chief Scientist of @WalmartLabs, a newly established research and development lab of Walmart.



Michael J. Franklin is a Professor of Computer Science at the UC, Berkeley, where he is a Director of the newly-opened Algorithms, Machines and People Laboratory (AMPLab). AMPLab is a five-year, industry- supported collaboration working at the intersection of statistical machine learning (Algorithms), cloud computing (Machines), and crowdsourcing (People). AMPLab aims to combine these resources into a new generation of analytics systems for making sense of diverse data at scale. The CrowdDB hybrid crowd/cloud query processing system is an early effort in this direction.



Donald Kossmann is a professor for Computer Science at ETH Zurich (Switzerland). He completed his PhD at the Technical University of Aachen. After that, he held positions at the University of Maryland, the IBM Almaden Research Center, the University of Passau, the Technical University of Munich, and the University of Heidelberg. He is an ACM fellow, member of the board of trustees of the VLDB endowment, and was the PC chair of the ACM SIGMOD Conf., 2009. He is a co-founder of i-TV-T (1998), XQRL Inc., and 28msec Inc.(2006). His research interests lie in the area of databases and information systems.

Tim Kraska is a PostDoc in the AMP Lab at the Computer Science Division

of the UC, Berkeley. Before joining UC Berkeley, he received his PhD from

ETH Zurich. Currently, his research focuses on the intersection between

data management for analytics, machine learning, and crowdsourcing. As

cloud data management system to answer gueries that neither database

systems nor search engines can adequately answer.



VLDB 2011

Tutorial 5: Graph Data Management Systems for New Application Domains

Philippe Cudré-Mauroux and Sameh Elnikety

Date: Thursday, September 1, 2011 Time: 10:30 AM-12:00 PM Room: Cascade 1BC

Abstract: Graph data management has long been a topic of interest for database researchers. The topic gained renewed interest recently, motivated by the rapid emergence of new application domains including social networks and the Web of data. This tutorial characterizes graph data management techniques and categorizes recent graph data management systems. In this context, we focus on the management of very large graphs such as social networks or the Web of data, rather than on the management of many smaller graphs (which frequently appear in bioinformatics and cheminformatics). The first part of this tutorial describes the requirements imposed by new application domains, and provides a classification of recent systems according to their data and computation models. Our classification also highlights the main representations used to store the graph (dense/sparse native graphs, triple storage or relational layouts), and the access patterns and typical queries considered (reachability or neighborhood queries, updates versus reads, transactional requirements and graph consistency models). In the second part of this tutorial, we map the data and computation models to concrete graph management systems, highlighting target application domains, implementation techniques, scalability and workload requirements.

Philippe Cudré-Mauroux is an Associate NSF Professor at the University of Fribourg in Switzerland. Previously, he was a postdoctoral associate working in the database systems group at MIT. He received his Ph.D. from the Swiss Federal Institute of Technology in Lausanne (EPFL), where he won both the Doctorate Award and the EPFL Press Mention in 2007. Before joining the University of Fribourg, he worked on distributed information management systems for HP, IBM T.J. Watson Research, and Microsoft Research Asia. His research interests are in exascale data management and infrastructures for non-relational data. He was the main investigator of the GridVine decentralized RDF storage system and is currently building



dipLODocus[RDF], a scalable and efficient back-end to store and analyze very large graphs of Web data in the cloud. He will be PC Chair of the International Semantic Web conference in 2012 in Boston.

Sameh Elnikety is a researcher at Microsoft Research in Redmond, Washington. He received his Ph.D. from the Swiss Federal Institute of Technology (EPFL) in Lausanne, Switzerland, and a M.Sc. from Rice University in Houston, Texas. His research interests include social network systems, graph databases, and large-scale software systems. Sameh is currently building Horton, a distributed system than manages large graphs on commodity servers while providing a declarative query language with multi-version transactional support. Sameh Elnikety is the PC Chair of Social Network Systems (SNS 2011) workshop, and the keynote speaker at Graph Data Management (GDM 2011) workshop.



Sameh's work on database replication received the best paper award at Eurosys 2007.

Tutorial 6: Information Diffusion In Social Networks: Observing and Influencing Societal Interests

Divyakant Agrawal, Ceren Budak, and Amr El Abbadi

Date: Thursday, September 1, 2011 Time: 1:30 PM-6:00 PM Room: Cascade 1BC

Abstract:: Social networks provide great opportunities for social connection, learning, political and social change, as well as individual entertainment and enhancement in a wide variety of forms. Because many social interactions currently take place in online networks, social scientists have access to unprecedented amounts of information about social interaction. This wealth of data can allow scientists to study social interactions on a scale and at a level of detail that has never been possible before.

In addition to providing a platform for scientists to observe social interactions at large scale, online social networks are also changing the very nature of social interactions. Understanding information diffusion in social networks is a critical research goal. Greater understanding can be achieved through data analysis, the development of reliable models that can predict outcomes of social processes, and ultimately the creation of applications that can shape the outcome of these processes. In this tutorial, we aim to provide an overview of such recent research based on a wide variety of techniques such as optimization algorithms, data mining, data streams covering a large number of problems such as influence spread maximization, misinformation limitation and study of trends in online social networks.

Divyakant Agrawal is a Professor of the Department of Computer Science at University of California, Santa Barbara. Prof. Agrawal's research expertise is in the areas of database systems, distributed computing, data warehousing, and large-scale information systems.



Ceren Budak is a PhD Candidate at the Department of Computer Science, University of California Santa Barbara. Her research interests lie in the area of information diffusion in social networks.



Amr El Abbadi is a Professor of the Department of Computer Science at University of California, Santa Barbara. His research interests lie in the area of scalable database and distributed systems.



Panel 1: Data Management for Meeting Global Health Challenges

Date: Tuesday, August 30, 2011 Time: 4:00PM-6:00PM Room: Fifth Avenue Authors: Tapan S. Parikh (UC Berkeley), Kuang Chen (UC Berkeley) Moderator: Tapan S. Parikh (UC Berkeley) Panelists: Kuang Chen (UC Berkeley) Lucky Gunasekara (Global Viral Forecasting Initiative) Alon Halevy (Google Inc.) Andy Kanter (Columbia University) Rowena Luk (Dimagi Inc.) Peter Speyer (University of Washington)

Abstract: Organizations and health agencies working on meeting global health challenges are becoming increasingly data driven. Governments and donors are demanding rigorous and up-to-date assessment of patient needs and health service delivery. These organizations work in some of the most difficult data environments - dealing with limited infrastructure and human resources, inconsistent data quality and availability, not to mention significant data backlogs.

This panel combines perspectives from a range of active researchers and practitioners managing large-scale deployments, including: evaluating large-scale health policies and interventions, managing day-to-day clinical activities, and responding to remote outbreaks of disease.

Data challenges in global health intersect with many current topics in database research. Potential topics to be discussed include: trade-offs between data quality, latency and cost; dealing with heterogeneous data sources, privacy challenges, managing provenance and mobile data collection.

Moderator: Tapan S. Parikh (UC Berkeley)

Tapan Parikh is an Assistant Professor at the School of Information at the University of

California, Berkeley. Tapan's research interests include humancomputer interaction (HCI), mobile computing, speech UIs and information systems for microfinance, smallholder agriculture and global health. For the past ten years, Tapan has been designing, developing and deploying information systems in the rural developing world - initially in India, and now also in Latin America and Africa. Tapan and his students have started several technology companies serving rural communities and the development sector. He holds a Sc.B. degree in Molecular Modeling with Honors from Brown University, and M.S. and Ph.D. degrees in Computer Science from the University of Washington, where he won the William Chan award for his Ph.D. dissertation. Tapan was named Technology Review magazine's



Humanitarian of the Year in 2007, for his work designing accessible mobile services for microfinance groups in rural India.

Panel 2: Panel Discussion: Maximizing Impact

Date: Wednesday, August 31, 2011 Time: 3:30PM-5:30PM Room: Vashon Authors: Ed Lazowska (University of Washington) Moderator: Ed Lazowska (University of Washington) Panelists: David DeWitt (Microsoft) Juliana Freire (NYU Polytechnic) Ed Lazowska (University of Washington) Sam Madden (MIT) Jennifer Widom (Stanford)

Abstract: This session is intended as an open discussion between a small group of panelists and the conference attendees. What can we do as a community to increase the impact (whatever that means) of our research, especially with the rise in paper counts, write only conferences, over-specialization of ourselves and our students, migration of distinguished researchers to more development-oriented positions, etc.? More specifically: What should we look for in graduate applicants? What should be the balance between breadth and depth in graduate programs? Are there emerging high-impact areas on which there should be greater focus (e.g., visualization, scientific data management)? How do we make data management techniques (databases, data mining, machine learning) usable by the masses? How can we contribute to a data-centric view of everything science, engineering, the social sciences, commerce?

Moderator: Ed Lazowska (University of Washington)

Ed Lazowska is the Bill & Melinda Gates Chair in Computer Science & Engineering at the

University of Washington, and is the founding Director of the University of Washington eScience Institute. He serves as Chair of the Computing Community Consortium, a national effort to engage the computing research community in fundamental research motivated by tackling societal challenges. He is a Member of the National Academy of Engineering, a Fellow of the American Academy of Arts & Sciences, and a Fellow of ACM, IEEE, and AAAS.



DEMONSTRATIONS

Demo A: Information Integration and Information Retrieval

Location: Grand Crescent

Tuesday 10:30 AM-12:00 PM; Wednesday 3:30-5:30 PM; Thursday 3:30-5:30 PM

CerFix: A System for Cleaning Data with Certain Fixes

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

Debugging Data Exchange with Vagabond

Boris Glavic (University of Toronto), Jiang Du (University of Toronto), Renee J. Miller (University of Toronto), Gustavo Alonso (ETH Zurich) Laura M. Haas (IBM Research - Almaden)

HOMES: A Higher-Order Mapping Evaluation System

Huy Vu (Oxford University), Michael Benedikt (Oxford University)

EIRENE: Interactive Design and Refinement of Schema Mappings via Data Examples

Bogdan Alexe (UC Santa Cruz), Balder ten Cate (UCSC), Phokion Kolaitis (UCSC & IBM Research - Almaden), Wang-Chiew Tan (IBM Research - Almaden & UCSC)

++Spicy: an Open-Source Tool for Second-Generation Schema Mapping and Data Exchange

Bruno Marnette (INRIA Saclay & ENS Cachan), Giansalvatore Mecca (Università della Basilicata), Paolo Papotti (Università Roma Tre), Salvatore Raunich (University of Leipzig), Donatello Santoro (Università della Basilicata)

Microsoft Codename "Montego" - Data Import, Transformation, and Publication for Information Workers

Stephen Maine (Microsoft Corporation), Lorenz Prem (Microsoft Corporation), Clemens Szyperski (Microsoft Corporation), James Terwilliger (Microsoft Corporation)

BROAD: Diversified Keyword Search in Databases

Feng Zhao (National University of Singapore), Xiaolong Zhang (Zhejiang University), Anthony Tung (National University of Singapore), Gang Chen (Zhejiang University)

DivDB: A System for Diversifying Query Results

Marcos Vieira (UCR), Humberto Razente (UFABC), Maria Camila Barioni (UFABC), Marios Hadjieleftheriou (AT&T Labs - Research), Divesh Srivastava (AT&T Labs), Caetano Traina Jr. (ICMC-USP), Vassilis Tsotras (UCR)

FuDoCS: A Web Service Composition System Based on Fuzzy Dominance for Preference Query Answering

Karim Benouaret (University of Lyon), Djamal Benslimane (University of Lyon), Allel Hadjali (University of Rennes), Mahmoud Barhamgi (University of Lyon)

AIDA: An Online Tool for Accurate Disambiguation of Named Entities in Text and Tables Mohamed Amir Yosef (Max-Planck-Institut für Informatik), Johannes Hoffart (Max-Planck-Institut für Informatik), Ilaria Bordino (Yahoo! Research), Marc Spaniol (Max-Planck-Institut für Informatik), Gerhard Weikum (Max-Planck-Institut für Informatik)

Demo B: Modern Hardware, Streaming, and Benchmarking

Location: Grand Crescent

Tuesday 1:30-3:30 PM; Wednesday 10:30 AM-12:00 PM; Thursday 1:30-3:00 PM

TrustedDB: A Trusted Hardware based Outsourced Database Engine Sumeet Bajaj (Stony Brook University), Radu Sion (Stony Brook University)

IPL-P: In-Page Logging with PCRAM

Kang-Nyeon Kim (Sungkyunkwan University), Sang-Won Lee (Sungkyunkwan University), Bongki Moon (University of Arizona), Chanik Park (Samsung Electronics), Joo-Young Hwang (Samsung Electronics Co., Ltd.)

HyPer-sonic Combined Transaction AND Query Processing

Florian Funke (Technische Universität München), Alfons Kemper (Technische Universität München), Thomas Neumann (Technische Universität München)

Proactive Detection and Repair of Data Corruption: Towards a Hassle-free Declarative Approach with Amulet

Nedyalko Borisov (Duke University), Shivnath Babu (Duke University)

A Demonstration of HYRISE - A Main Memory Hybrid Storage Engine

Martin Grund (Hasso-Plattner-Institut), Philippe Cudre-Mauroux (MIT CSAIL), Samuel Madden (MIT)

Analytics for the Real-Time Web

Maxim Grinev (ETH Zurich), Maria Grineva (ETH Zurich), Martin Hentschel (ETH Zurich), Donald Kossmann (ETH Zurich)

UpStream: A Storage-centric Load Management System for Real-time Update Streams Alexandru Moga (ETH Zurich), Nesime Tatbul (ETH Zurich)

DataSynth: Generating Synthetic Data using Declarative Constraints Arvind Arasu (Microsoft Research), Raghav Kaushik (Microsoft Research),

Jian Li (University of Maryland)

Automatic Workload Driven Index Defragmentation

Vivek Narasayya (Microsoft Research), Hyunjung Park (Stanford University), Manoj Syamala (Microsoft Research)

Demo C: MapReduce, Crowdsourcing, and Mining

Location: Grand Crescent

Tuesday 4:00-6:00 PM; Wednesday 1:30-3:00 PM; Thursday 10:30 AM-12:00 PM

Whodunit: An Auditing Tool for Detecting Data Breaches

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

CrowdDB: Query Processing with the VLDB Crowd

Amber Feng (UC Berkeley), Michael Franklin (UC Berkeley), Donald Kossmann

(ETH Zurich), Tim Kraska (UC Berkeley), Samuel Madden (MIT), Sukriti Ramesh (Systems Group, ETH Zurich), Andrew Wang (UC Berkeley), Reynold Xin (UC Berkeley)

InfoNetOLAPer: Integrating InfoNetWarehouse and InfoNetCube with InfoNetOLAP

Chuan Li (UIC), Philip Yu (University of Illinois at Chicago), Lei Zhao (University of Science and Technology of China), Yan Xie (University of Illinois at Chicago), Wangqun Lin (University of Illinois at Chicago)

RAMP: A System for Capturing and Tracing Provenance in MapReduce Workflows Hyunjung Park (Stanford University), Robert Ikeda (Stanford University), Jennifer Widom (Stanford University)

From SPARQL to MapReduce: The Journey Using a Nested TripleGroup Algebra

HyeongSik Kim (North Carolina State University), Padmashree Ravindra (North Carolina State University), Kemafor Anyanwu (North Carolina State University)

MapReduce Programming and Cost-based Optimization? Crossing this Chasm with Starfish

Herodotos Herodotou (Duke University), Fei Dong (Duke University), Shivnath Babu (Duke University)

SocialSpamGuard: A Data Mining-Based Spam Detection System for Social Media Networks

Xin Jin (UIUC), Cindy Xide Lin (UIUC), Jiebo Luo (Kodak Research), Jiawei Han (UIUC)

GrouPeer: A System for Clustering PDMSs

Verena Kantere (Cypress University of Technology), Dimos Bousounis (ETH Zurich), Timos Sellis (National Technical University of Athens)

Online Visualization of Geospatial Stream Data using the WorldWide Telescope Mohamed Ali (Microsoft), Badrish Chandramouli (Microsoft Research), Jonathan Fay (Microsoft), Curtis Wong (Microsoft), Steven Drucker (Microsoft), Balan Sethu Raman (Microsoft)

Research 1: Distributed Systems

Distributed Threshold Querying of General Functions by a Difference of Monotonic Representation

Guy Sagy (Technion), Daniel Keren (Haifa University), Izchak Sharfman (Technion), Assaf Schuster (Technion)

The goal of a threshold query is to detect all objects whose score exceeds a given threshold. This type of query is used in many settings, such as data mining, event triggering, and top-k selection. Often, threshold queries are performed over distributed data. Given database relations that are distributed over many nodes, an object's score is computed by aggregating the value of each attribute, applying a given scoring function over the aggregation, and thresholding the function's value. However, joining all the distributed relations to a central database might incur prohibitive overheads in bandwidth, CPU, and storage accesses. Efficient algorithms required to reduce these costs exist only for monotonic aggregation threshold queries and certain specific scoring functions. We present a novel approach for efficiently performing general distributed threshold queries. To the best of our knowledge, this is the first general solution to the problem of performing such queries with arbitrary scoring functions. We first present a solution for monotonic functions, and then introduce a technique to solve for other functions by representing them as a difference of monotonic functions. Experiments with real-world data demonstrate the method's effectiveness in achieving low communication and access costs.

Distributed Inference and Query Processing for RFID Tracking and Monitoring

Zhao Cao (University of Massachusetts), Charles Sutton (University of Edinburgh), Yanlei Diao (University of Massachusetts), Prashant Shenoy (University of Massachusetts)

In this paper, we present the design of a scalable, distributed stream processing system for RFID tracking and monitoring. Since RFID data lacks containment and location information that is key to query processing, we propose to combine location and containment inference with stream query processing in a single architecture, with inference as an enabling mechanism for high-level query processing. We further consider challenges in instantiating such a system in large distributed settings and design techniques for distributed inference and query processing. Our experimental results, using both real-world data and large synthetic traces, demonstrate the accuracy, efficiency, and scalability of our proposed techniques.

Where in the World is My Data?

Sudarshan Kadambi (Bloomberg), Jianjun Chen (Yahoo!), Brian Cooper (Google), David Lomax (Yahoo!), Raghu Ramakrishnan (Yahoo!), Adam Silberstein (Yahoo!), Erwin Tam (Yahoo!), Hector Garcia-Molina (Stanford University)

Users of websites such as Facebook, Ebay and Yahoo! demand fast response times, and these sites replicate data across globally distributed datacenters to achieve this. However, it is not necessary to replicate all data to all locations: if a European user's record is never accessed in Asia, it does not make sense to pay the bandwidth and disk costs to maintain an Asian replica. In this paper, we describe mechanisms for selectively replicating large-scale web databases on a record-by-record basis. We introduce a flexible constraint language to specify replication policy constraints. We then present an adaptive scheme for replicating data to where it is most frequently accessed, while respecting policy constraints and using minimal bookkeeping. Experiments using a modified version of our PNUTS system demonstrate our techniques work well.

Research 2: Entity Matching

Entity Matching: How Similar Is Similar?

Jiannan Wang (Tsinghua University), Guoliang Li (Tsinghua University), Jeffrey Xu Yu (Chinese University of Hong Kong), Jianhua Feng (Tsinghua University) Entity matching that finds records referring to the same entity is an important operation in data cleaning and integration. Existing studies usually use a given similarity function to quantify the similarity of records, and focus on devising index structures and algorithms for efficient entity matching. However it is a big challenge to define "how similar is similar" for real applications, since it is rather hard to automatically select appropriate similarity functions. In this paper we attempt to address this problem. As there are a large number of similarity functions, and even worse thresholds may have infinite values, it is rather expensive to find appropriate similarity functions and thresholds. Fortunately, we have an observation that different similarity functions and thresholds have redundancy, and we have an opportunity to prune inappropriate similarity functions. To this end, we propose effective optimization techniques to eliminate such redundancy, and devise efficient algorithms to find the best similarity functions. The experimental results on both real and synthetic datasets show that our method achieves high accuracy and outperforms the baseline algorithms.

Large-Scale Collective Entity Matching

Vibhor Rastogi (Yahoo! Research), Nilesh Dalvi (Yahoo! Research), Minos Garofalakis (Technical University of Crete)

There have been several recent advancements in Machine Learning community on the Entity Matching (EM) problem. However, their lack of scalability has prevented them from being applied in practical settings on large real-life datasets. Towards this end, we propose a principled framework to scale any generic EM algorithm. Our technique consists of running multiple instances of the EM algorithm on small neighborhoods of the data and passing messages across neighborhoods to construct a global solution. We prove formal properties of our framework and experimentally demonstrate the effectiveness of our approach in scaling EM algorithms.

Linking Temporal Records

Pei Li (University of Milan - Bicocca), Xin Dong (AT&T Labs), Andrea Maurino (University of Milan - Bicocca), Divesh Srivastava (AT&T Labs)

Many data sets contain temporal records over a long period of time; each record is associated with a time stamp and describes some aspects of a real-world entity at that particular time (e.g., author information in DBLP). In such cases, we often wish to identify records that describe the same entity over time and so be able to enable interesting longitudinal data analysis. However, existing record linkage techniques ignore the temporal information and can fall short for temporal data. This paper studies linking temporal records. First, we apply time decay to capture the effect of elapsed time on entity value evolution. Second, instead of comparing each pair of records locally, we propose clustering methods that consider time order of the records and make global decisions. Experimental results show that our algorithms significantly outperform traditional linkage methods on various temporal data sets.

Research 3: Web

Output URL Bidding

Panagiotis Papadimitriou (Stanford University), Hector Garcia-Molina (Stanford University), Ali Dasdan (Ebay Inc), Santanu Kolay (Ebay Inc)

Output URL bidding is a new bidding mechanism for sponsored search, where advertisers bid on search result URLs, as opposed to keywords in the input query. For example, an advertiser may want his ad to appear whenever the search result includes the sites www.imdb.com and en.wikipedia.org, instead of bidding on keywords that lead to these sites, e.g., movie titles or actor names. In this paper we study the tradeoff between the simplicity and the specification power of output bids and we explore their utility for advertisers. We first present a model to derive output bids from existing keyword bids. Then, we use the derived bids to experimentally

study output bids and contrast them to input query bids. Our main results are the following: (1) Compact output bids that mix both URLs and hosts have the same specification power as more lengthy input bids; (2) Output bidding can increase the recall of relevant queries; and (3) Output and input biding can be combined into a hybrid mechanism that combines the benefits of both.

Automatic Wrappers for Large Scale Web Extraction

Nilesh Dalvi (Yahoo! Research), Ravi Kumar (Yahoo! Research), Mohamed Soliman (University of Waterloo)

We present a generic framework to make wrapper induction algorithms tolerant to noise in the training data. This enables us to learn wrappers in a completely unsupervised manner from automatically and cheaply obtained noisy training data, e.g., using dictionaries and regular expressions. By removing the site-level supervision that wrapper-based techniques require, we are able to perform information extraction at web-scale, with accuracy unattained with existing unsupervised extraction techniques. Our system is used in production at Yahoo! and powers live applications.

Recovering Semantics of Tables on the Web

Petros Venetis (Stanford University), Alon Halevy (Google), Jayant Madhavan (Google), Marius Pasca (Google), Warren Shen (Google), Fei Wu (Google), Gengxin Miao (UC Santa Barbara), Chung Wu (Google)

The Web offers a corpus of over 100 million tables, but the meaning of each table is rarely explicit from the table itself. Header rows exist in few cases and even when they do, the attribute names are typically useless. We describe a system that attempts to recover the semantics of tables by enriching the table with additional annotations. Our annotations facilitate operations such as searching for tables and finding related tables. To recover semantics of tables, we leverage a database of class labels and relationships automatically extracted from the Web. The database of classes and relationships has very wide coverage, but is also noisy. We attach a class label to a column if a sufficient number of the values in the column are identified with that label in the database of class labels, and analogously for binary relationships. We describe a formal model for reasoning about when we have seen sufficient evidence for a label, and show that it performs substantially better than a simple majority scheme. We describe a set of experiments that illustrate the utility of the recovered semantics for table search and show that it performs substantially better than previous approaches. In addition, we characterize what fraction of tables on the Web can be annotated using our approach.

Research 4: GeoSpatial

Graph Indexing of Road Networks for Shortest Path Queries with Label Restrictions Michael Rice (UCR), Vassilis Tsotras (UCR)

The current widespread use of location-based services and GPS technologies has revived interest in very fast and scalable shortest path queries. We introduce a new shortest path query type in which dynamic constraints may be placed on the allowable set of edges that can appear on a valid shortest path (e.g., dynamically restricting the type of roads or modes of travel which may be considered in a multimodal transportation network). We formalize this problem as a specific variant of formal language constrained shortest path problems, which we call the Kleene Language Constrained Shortest Paths problem. To efficiently support this type of dynamically constrained shortest path query for large-scale datasets, we extend the hierarchical graph indexing technique known as Contraction Hierarchies. Our experimental evaluation using the North American road network dataset (with over 50 million edges) shows an average query speed and search space improvement of over 3 orders of magnitude compared to the naive adaptation of the standard Dijkstra's algorithm to support this query type. We also show an improvement of over 2 orders of magnitude compared to the only previously-existing indexing technique which could solve this problem without additional preprocessing.

Efficient Processing of Top-k Spatial Preference Queries

João B. Rocha-Junior (NTNU), Akrivi Vlachou (NTNU), Christos Doulkeridis (NTNU), Kjetil Nørvåg (NTNU)

Top-k spatial preference queries return a ranked set of the k best data objects based on the scores of feature objects in their spatial neighborhood. Despite the wide range of location-based applications that rely on spatial preference queries, existing algorithms incur non-negligible processing cost resulting in high response time. The reason is that computing the score of any data object requires examining its spatial neighborhood to find the feature object with highest score. In this paper, we propose a novel technique to speed up the performance of top-k spatial preference queries. To this end, we propose a mapping of pairs of data and feature objects to a distance-score space, which in turn allows us to identify and materialize the minimal subset of pairs that is sufficient to answer a spatial preference query. Furthermore, we present a novel algorithm that improves query processing performance by avoiding examining the spatial neighborhood of the data objects during query execution. In addition, we propose an efficient algorithm for materialization and we describe useful properties that reduce the cost of maintenance. We show through extensive experiments that our approach significantly reduces the number of I/Os and execution time compared to the state-of-the-art algorithms for different setups.

SXPath - Extending XPath towards Spatial Querying on Web Documents Ermelinda Oro (DEIS-UNICAL, Altilia srl), Massimo Ruffolo (ICAR-CNR, Altilia srl), Steffen Staab (Institute WeST, University of Koblenz-Landau)

Querying data from presentation formats like HTML, for purposes such as information extraction, requires the consideration of tree structures as well as the consideration of spatial relationships between laid out elements. The underlying rationale is that frequently the rendering of tree structures is very involved and undergoing more frequent updates than the resulting layout structure. Therefore, in this paper, we present Spatial XPath (SXPath), an extension of XPath 1.0 that allows for inclusion of spatial navigation primitives into the language resulting in conceptually simpler queries on Web documents. The SXPath language is based on a combination of a spatial algebra with formal descriptions of XPath navigation, and maintains polynomial time combined complexity. Practical experiments demonstrate the usability of SXPath.

Efficient Algorithms for Finding Optimal Meeting Point on Road Networks Da Yan (HKUST), Zhou Zhao (HKUST), Wilfred Ng (HKUST)

Given a set of points Q on a road network, an optimal meeting point (OMP) query returns the point on a road network G=(V, E) with the smallest sum of network distances to all the points in Q. This problem has many real world applications, such as minimizing the total travel cost for a group of people who want to find a location for gathering. While this problem has been well studied in the Euclidean space, the recently proposed state-of-the-art algorithm for solving this problem in the context of road networks is still not efficient. In this paper, we propose a new baseline algorithm for the OMP query, which reduces the search space from |Q||E| to |V|+|Q|. We also present two effective pruning techniques that further accelerate the baseline algorithm. Finally, in order to support spatial applications that involve large flow of queries and require fast response, an extremely efficient algorithm is proposed to find a high-quality near-optimal meeting point, which is orders of magnitude faster than the exact OMP algorithms. Extensive experiments are conducted to verify the efficiency of our algorithms.

Research 5: Uncertain Data

A Generic Framework for Handling Uncertain Data with Local Correlations Xiang Lian (HKUST), Lei Chen (HKUST)

Data uncertainty is ubiquitous in many real-world applications such as sensor/RFID data analysis. In this paper, we investigate uncertain data that exhibit local correlations, that is, each uncertain object is only locally correlated with a small subset of data, while being independent of others. We propose a generic framework for dealing with this kind of uncertain and locally correlated data, in which we investigate a classical spatial query, nearest neighbor query, on uncertain data with local correlations (namely LC-PNN). Most importantly, to enable fast LC-PNN query processing, we propose a novel filtering technique via offline pre-computations to reduce the query search space. We demonstrate through extensive experiments the efficiency and effectiveness of our approaches.

Efficient Probabilistic Reverse Nearest Neighbor Query Processing on Uncertain Data

Thomas Bernecker (Ludwig-Maximilians-University), Tobias Emrich (Ludwig-Maximilians-University), Hans-Peter Kriegel (Ludwig-Maximilians-University), Matthias Renz (Ludwig-Maximilians-University), Stefan Zankl (Ludwig-Maximilians-University), Andreas Züfle (Ludwig-Maximilians-University)

Given a query object q, a reverse nearest neighbor (RNN) query in a common certain database returns the objects having q as their nearest neighbor. A new challenge for databases is dealing with uncertain objects. In this paper we consider probabilistic reverse nearest neighbor (PRNN) queries, which return the uncertain objects having the query object as nearest neighbor with a sufficiently high probability. We propose an algorithm for efficiently answering PRNN queries using new pruning mechanisms taking distance dependencies into account. We compare our algorithm to state-of-the-art approaches recently proposed. Our experimental evaluation shows that our approach is able to significantly outperform previous approaches. In addition, we show how our approach can easily be extended to PRkNN (where k > 1) query processing for which there is currently no efficient solution.

Queries with Difference on Probabilistic Databases Sanjeev Khanna (University of Pennsylvania), Sudeepa Roy (University of

Pennsylvania), Val Tannen (University of Pennsylvania)

We study the feasibility of the exact and approximate computation of the probability of relational queries with difference on tuple-independent databases. We show that even the difference between two "safe" conjunctive queries without self-joins is "unsafe" for exact computation. We turn to approximation and design an FPRAS for a large class of relational queries with difference, limited by how difference is nested and by the nature of the subtracted sub-queries. We give examples of inapproximable queries outside this class.

Optimizing Probabilistic Query Processing on Continuous Uncertain Data Liping Peng (UMass Amherst), Yanlei Diao (UMass Amherst), Anna Liu (UMass Amherst)

Uncertain data management is becoming increasingly important in many applications, in particular, in scientific databases and data stream systems. Uncertain data in these new environments is naturally modeled by continuous random variables. An important class of queries uses complex selection and join predicates and requires query answers to be returned if their existence probabilities pass a threshold. In this work, we optimize threshold query processing for continuous uncertain data by (i) expediting joins using new indexes on uncertain data, (ii) expediting selections by reducing dimensionality of integration and using faster filters, and (iii) optimizing a query plan using a dynamic, per-tuple based approach. Evaluation results using real-world data and benchmark queries show the accuracy and efficiency of our techniques and significant performance gains over a state-of-the-art threshold query optimizer.

Research 6: Database Design

CRIUS: User-Friendly Database Design

Li Qian (University of Michigan), Kristen LeFevre (University of Michigan), H. Jagadish (University of Michigan)

Non-technical users are increasingly adding structures to their data. This gives rise to the need for database design. However, traditional database design is deliberate and heavy-weight, requiring technical expertise that everyday users may not possess. For this reason, we propose that users of personal data management applications should be able to create and refine data structures in an ad-hoc way over time, thereby "organically" growing their schemas. For this purpose, we develop a spreadsheet-like direct manipulation interface. We show how integrity constraints can still provide value, even in this scenario of frequent schema and data modifications. We also develop a back-end database implementation to support this interface, with a design that permits schema changes at a low cost. We have folded these ideas into a system, called CRIUS, which supports a nested data model and a graphical user interface. From the user's perspective, the chief advantages of CRIUS are its support for simple schema definition and modification through an intuitive drag-and-drop interface, as well as its guidance towards user data entry based on incrementally updated data integrity. We have evaluated CRIUS by means of user studies and performance studies. The user studies indicate that 1) CRIUS makes it much easier for users to design a database, as compared to state-of-the-art GUI database design tools, and 2) CRIUS makes user data entry more efficient and less error-prone. The performance experiments show that 1) the incremental integrity update in CRIUS is very efficient, making the data entry guidance applicable and 2) the back-end database implementation in CRIUS significantly improves the performance of schema update tasks, without a significant impact on other operations.

CoPhy: A Scalable, Portable, and Interactive Index Advisor for Large Workloads Debabrata Dash (ArcSight), Neoklis Polyzotis (UC Santa Cruz), Anastasia Ailamaki (EPFL)

Index tuning, i.e., selecting the indexes appropriate for a workload, is a crucial problem in database system tuning. In this paper, we solve index tuning for large problem instances that are common in practice, e.g., thousands of queries in the workload, thousands of candidate indexes and several hard and soft constraints. Our work is the first to reveal that the index tuning problem has a well structured space of solutions, and this space can be explored efficiently with well known techniques from linear optimization. Experimental results demonstrate that our approach outperforms state-of-the-art commercial and research techniques by a significant margin (up to an order of magnitude).

Compression Aware Physical Database Design Hideaki Kimura (Brown University), Vivek Narasayya (Microsoft Research), Manoj Syamala (Microsoft Research)

Modern RDBMSs support the ability to compress data using methods such as null suppression and dictionary encoding. Data compression offers the promise of significantly reducing storage requirements and improving I/O performance for decision support queries. However, compression can also slow down update and query performance due to the CPU costs of compression and decompression. In this paper, we study how data compression affects choice of appropriate physical database design, such as indexes, for a given workload. We observe that approaches that decouple the decision of whether or not to choose an index from whether or not to compress the index can result in poor solutions. Thus, we focus on the novel problem of integrating compression into physical database design in a scalable manner. We have implemented our techniques by modifying Microsoft SQL Server and the Database Engine Tuning

Advisor (DTA) physical design tool. Our techniques are general and are potentially applicable to DBMSs that support other compression methods. Our experimental results on real world as well as TPC-H benchmark workloads demonstrate the effectiveness of our techniques.

Summary Graphs for Relational Database Schemas

Xiaoyan Yang (National University of Singapore), Cecilia Procopiuc (AT&T Labs), Divesh Srivastava (AT&T Labs)

Increasingly complex databases require ever more sophisticated tools to help users understand their schemas and interact with the data. Existing tools fall short of either providing the "big picture", or of presenting useful connectivity information. In this paper we define summary graphs, a novel approach for summarizing schemas. Given a set of user-specified query tables, the summary graph automatically computes the most relevant tables and joins for that query set. The output preserves the most informative join paths between the query tables, while meeting size constraints. In the process, we define a novel information-theoretic measure over join edges. Unlike most subgraph extraction work, we allow metaedges (i.e., edges in the transitive closure) to help reduce output complexity. We prove that the problem is NP-Hard, and solve it as an integer program. Our extensive experimental study shows that our method returns high-quality summaries under independent quality measures.

Research 7: Query Processing

Generating Efficient Execution Plans for Vertically Partitioned XML Databases Patrick Kling (University of Waterloo), M. Tamer Özsu (University of Waterloo), Khuzaima Daudjee (University of Waterloo)

Experience with relational systems has shown that distribution is an effective way of improving the scalability of query evaluation. In this paper, we show how distributed query evaluation can be performed in a vertically partitioned XML database system. We propose a novel technique for constructing distributed execution plans that is independent of local query evaluation strategies. We then present a number of optimizations that allow us to further improve the performance of distributed query execution. Finally, we present a response time-based cost model that allows us to pick the best execution plan for a given query and database instance. Based on an implementation of our techniques within a native XML database system, we verify that our execution plans take advantage of the parallelism in a distributed system and that our cost model is effective at identifying the most advantageous plans.

Similarity Join Size Estimation using Locality Sensitive Hashing

Hongrae Lee (University of British Columbia), Raymond Ng (University of British Columbia), Kyuseok Shim (Seoul National University)

Similarity joins are important operations with a broad range of applications. In this paper, we study the problem of vector similarity join size estimation (VSJ). It is a generalization of the previously studied set similarity join size estimation (SSJ) problem and can handle more interesting cases such as TF-IDF vectors. One of the key challenges in similarity join size estimation is that the join size can change dramatically depending on the input similarity threshold. We propose a sampling based algorithm that uses the Locality-Sensitive-Hashing (LSH) scheme. The proposed algorithm LSH-SS uses an LSH index to enable effective sampling even at high thresholds. We compare the proposed technique with random sampling and the state-of-the-art technique for SSJ (adapted to VSJ) and demonstrate LSH-SS offers more accurate estimates at both high and low similarity thresholds and small variance using real-world data sets.

Accelerating Queries with Group-By and Join by Groupjoin

Guido Moerkotte (University of Mannheim), Thomas Neumann (Technische Universität München)

Most aggregation queries contain both group-by and join operators, and spend a significant amount of time evaluating these two expensive operators. Merging them into one operator (the groupjoin) significantly speeds up query execution. We introduce two main equivalences to allow for the merging and prove their correctness. Furthermore, we show experimentally that these equivalences can significantly speed up TPC-H.

Optimizing Query Answering under Ontological Constraints Giorgio Orsi (Oxford University), Andreas Pieris (Oxford University)

Ontological queries are evaluated against a database combined with ontological constraints and answering such queries is a challenging new problem for database research. For many ontological modelling languages, query answering can be solved via query rewriting: given a conjunctive query and an ontology, the query can be transformed into a first-order query, called the perfect rewriting, that takes into account the semantic consequences of the ontology. Then, for every extensional database D, the answer to the query is obtained by evaluating the rewritten query against D. In this paper we present a new algorithm that computes the perfect rewriting of a conjunctive query w.r.t. a linear Datalog+- ontology. Also, we provide an experimental comparison of our algorithm with existing rewriting techniques.

Research 8: Graph Data

On Triangulation-based Dense Neighborhood Graph Discovery

Nan Wang (National University of Singapore), Jingbo Zhang (National University of Singapore), Kian-Lee Tan (National University of Singapore), Anthony Tung (National University of Singapore)

This paper introduces a new definition of dense subgraph pattern, the DN-graph. DN-graph considers both the size of the sub-structure and the minimum level of interactions between any pair of the vertices. The mining of DN-graphs inherits the difficulty of finding clique, the fully-connected subgraphs. We thus opt for approximately locating the DN-graphs using the state-of-the-art graph triangulation methods. Our solution consists of a family of algorithms, each of which targets a different problem setting. These algorithms are iterative, and utilize repeated scans through the triangles in the graph to approximately locate the DN-graphs. Each scan on the graph triangles improves the results. Since the triangles are not physically materialized, the algorithms have small memory footprint. With our solution, the users can adopt a "pay as you go" approach. They have the flexibility to terminate the mining process once they are satisfied with the quality of the results. As a result, our algorithms can cope with semi-streaming environment where the graph edges cannot fit into main memory. Results of extensive performance study confirmed our claims.

Human-Assisted Graph Search: It's Okay to Ask Questions

Aditya Parameswaran (Stanford University), Anish Das Sarma (Yahoo! Research), Hector Garcia-Molina (Stanford University), Neoklis Polyzotis (UC Santa Cruz), Jennifer Widom (Stanford University)

We consider the problem of human-assisted graph search: given a directed acyclic graph with some (unknown) target node(s), we consider the problem of finding the target node(s) by asking an omniscient human questions of the form "Is there a target node that is reachable from the current node?". This general problem has applications in many domains that can utilize human intelligence, including curation of hierarchies, debugging workflows, image segmentation and categorization, interactive search and filter synthesis. To our knowledge, this work provides

the first formal algorithmic study of the optimization of human computation for this problem. We study various dimensions of the problem space, providing algorithms and complexity results. Our framework and algorithms can be used in the design of an optimizer for crowdsourcing platforms such as Mechanical Turk.

On Querying Historical Evolving Graph Sequences

Chenghui Ren (University of Hong Kong), Eric Lo (HK Polytechnic University), Ben Kao (University of Hong Kong), Xinjie Zhu (University of Hong Kong), Reynold Cheng (University of Hong Kong)

In many applications, information is best represented as graphs. In a dynamic world, information changes and so the graphs representing the information evolve with time. We propose that historical graph-structured data be maintained for analytical processing. We call a historical evolving graph sequence an EGS. We observe that in many applications, graphs of an EGS are large and numerous, and they often exhibit much redundancy among them. We study the problem of efficient query processing on an EGS and put forward a solution framework called FVF. Through extensive experiments on both real and synthetic datasets, we show that our FVF framework is highly efficient in EGS query processing.

Mining Top-K Large Structural Patterns in a Massive Network

Feida Zhu (Singapore Management University), Qiang Qu (Peking University), David Lo (Singapore Management University), Xifeng Yan (UCSB), Jiawei Han (UIUC), Philip Yu (UIC)

With ever-growing popularity of social networks, web and bio-networks, mining large frequent patterns from a single huge network has become increasingly important. Yet the existing pattern mining methods cannot offer the efficiency desirable for large pattern discovery. We propose SpiderMine, a novel algorithm to efficiently mine top-K largest frequent patterns from a single massive network with any user-specified probability of 1- ϵ . Deviating from the existing edge-by-edge (ie, incremental) pattern-growth framework, SpiderMine achieves its efficiency by unleashing the power of small patterns of a bounded diameter, which we call "spiders". With the spider structure, our approach adopts a probabilistic mining framework to find the top-k largest patterns by (i) identifying an affordable set of promising growth paths toward large patterns, and (ii) generating large patterns with much lower combinatorial complexity, and finally (iii) greatly reducing the cost of graph isomorphism tests with a new graph pattern representation by a multi-set of spiders. Extensive experimental studies on both synthetic and real data sets show that our algorithm outperforms existing methods.

Research 9: New Hardware Architecture

HYRISE - A Main Memory Hybrid Storage Engine

Martin Grund (Hasso-Plattner-Institute), Jens Krüger (Hasso-Plattner-Institute), Hasso Plattner (Hasso-Plattner Institute), Alexander Zeier (Hasso-Plattner Institute), Philippe Cudre-Mauroux (MIT), Samuel Madden (MIT)

In this paper, we describe a main memory hybrid database system called HYRISE, which automatically partitions tables into vertical partitions of varying widths depending on how the columns of the table are accessed. For columns accessed as a part of analytical queries (e.g., via sequential scans), narrow partitions perform better, because, when scanning a single column, cache locality is improved if the values of that column are stored contiguously. In contrast, for columns accessed as a part of OLTP-style queries, wider partitions perform better, because such transactions frequently insert, delete, update, or access many of the fields of a row, and co-locating those fields leads to better cache locality. Using a highly accurate model of cache misses, HYRISE is able to predict the performance of different partitionings, and to automatically select the best partitioning using an automated database design algorithm. We show that, on a realistic workload derived from customer applications, HYRISE can achieve a 20% to 400% performance improvement over pure all-column or all-row designs, and that it is both more scalable and produces better designs than previous vertical partitioning approaches for main memory systems.

Fast Set Intersection in Memory Bolin Ding (UIUC), Arnd Christian König (Microsoft Research)

Set intersection is a fundamental operation in information retrieval and database systems. This paper introduces linear space data structures to represent sets such that their intersection can be computed in a worst-case efficient way. In general, given k (preprocessed) sets, with totally n elements, we will show how to compute their intersection in expected time O(n/(sqrt(w))+kr), where r is the intersection size and w is the number of bits in a machine-word. In addition,we introduce a very simple version of this algorithm that has weaker asymptotic guarantees but performs even better in practice; both algorithms outperform the state of the art techniques in terms of execution time for both synthetic and real data sets and workloads.

Efficiently Compiling Efficient Query Plans for Modern Hardware Thomas Neumann (Technische Universität München)

As main memory grows, query performance is more and more determined by the raw CPU costs of query processing itself. The classical iterator style query processing technique is very simple and flexible, but shows poor performance on modern CPUs due to lack of locality and frequent instruction mis-predictions. Several techniques like batch oriented processing or vectorized tuple processing have been proposed in the past to improve this situation, but even these techniques are frequently out-performed by hand-written execution plans. In this work we present a novel compilation strategy that translates a query into compact and efficient machine code using the LLVM compiler framework. By aiming at good code and data locality and predictable branch layout the resulting code frequently rivals the performance of hand-written C++ code. We integrated these techniques into the HyPer main memory database system and show that this results in excellent query performance while requiring only modest compilation time.

PALM: Parallel Architecture-Friendly Latch-Free Modifications to B+ Trees on Many-Core Processors

Jason Sewall (Intel Corporation), Jatin Chhugani (Intel Corporation), Changkyu Kim (Intel Corporation), Nadathur Satish (Intel Corporation), Pradeep Dubey (Intel Corporation)

Concurrency control on B+ trees is primarily achieved with latches, but serialization and contention can hinder scalability. As core counts on current processors increase, it is imperative to develop scalable latch-free techniques for concurrency control. We present PALM, a novel technique for performing multiple concurrent queries on in-memory B+ trees. PALM is based on the Bulk Synchronous Parallel model, which guarantees freedom from deadlocks and race conditions. Input queries are grouped and processed in atomic batches, and work proceeds in stages that preclude contention. Transition between stages is accomplished with scalable pointto-point communication. PALM exploits data-and thread-level parallelism on modern manycore architectures, and performs 40M updates/second on trees with 128M keys, and 128M updates/second on trees with 512K keys on the latest CPU architectures. Our throughput is 2.3X-19X that of state-of-the-art concurrent update algorithms on in-memory B+ trees. PALM obtains close to peak throughput at very low response times of less than 350microseconds, even for large trees. We also evaluate PALM on the Intel(R) Many Integrated Core (Intel(R) MIC) architecture, and demonstrate a speedup of 1.5-2.1X for out-of-cache tree sizes on an Intel(R) Knights Ferry over a pair of Intel(R) Xeon(R) processors DP X5680 (Westmere-EP)in a dualsocket configuration.

Research 10: Causality, Quality, and Dependencies

The Complexity of Causality and Responsibility for Query Answers and non-Answers Alexandra Meliou (University of Washington), Wolfgang Gatterbauer (University of Washington), Katherine Moore (University of Washington), Dan Suciu (University of Washington)

An answer to a query has a well-defined lineage expression (alternatively called how-provenance) that explains how the answer was derived. Recent work has also shown how to compute the lineage of a non-answer to a query. However, the cause of an answer or non-answer is a more subtle notion and consists, in general, of only a fragment of the lineage. In this paper, we adapt Halpern, Pearl, and Chockler's recent definitions of causality and responsibility to define the causes and responsibilities of answers and non-answers to queries, and their degree of responsibility. Responsibility captures the notion of degree of causality and serves to rank potentially many causes by their relative contributions to the effect. Then, we study the complexity of computing causes and responsibilities for conjunctive queries. It is known that computing causes is in general NP-complete. Our first main result shows that all causes to conjunctive queries can be computed by a relational query which may involve negation. Thus, causality can be computed in PTIME, and very efficiently so. Next, we study computing responsibility. Here, we prove that the complexity depends on the conjunctive query and demonstrate a dichotomy between PTIME and NP-complete cases. For the PTIME cases we give a non-trivial algorithm, consisting of a reduction to the max-flow computation problem. Finally, we prove that, even when it is in PTIME, responsibility is complete for LOGSPACE, implying that, unlike causality, it cannot be computed by a relational query.

Guided Data Repair

Mohamed Yakout (Purdue University), Ahmed Elmagarmid (Qatar Computing Research Institute), Jennifer Neville (Purdue University), Mourad Ouzzani (Purdue University), Ihab Ilyas (University of Waterloo)

In this paper we present GDR, a Guided Data Repair framework that incorporates user feedback in the cleaning process to enhance and accelerate existing automatic repair techniques while minimizing user involvement. GDR consults the user on the updates that are most likely to be beneficial in improving data quality. GDR also uses machine learning methods to identify and apply the correct updates directly to the database without the actual involvement of the user on these specific updates. To rank potential updates for consultation by the user, we first group these repairs and quantify the utility of each group using the decision-theory concept of value of information (VOI). We then apply active learning to order updates within a group based on their ability to improve the learned model. User feedback is used to repair the database and to adaptively refine the training set for the model. We empirically evaluate GDR on a real-world dataset and show significant improvement in data quality using our user guided repairing process. We also, assess the trade-off between the user efforts and the resulting data quality.

Stratification Criteria and Rewriting Techniques for Checking Chase Termination Sergio Greco (Università della Calabria), Francesca Spezzano (Università della Calabria), Irina Trubitsyna (DEIS, Università della Calabria)

The Chase is a fixpoint algorithm enforcing satisfaction of data dependencies in databases. Its execution involves the insertion of tuples with possible null values and the changing of null values which can be made equal to constants or other null values. Since the chase fixpoint evaluation could be on-terminating, in recent years the problem know as chase termination has been investigated. It consists in the detection of sufficient conditions, derived from the structural analysis of dependencies, guaranteeing that the chase fixpoint terminates independencies.

dently from the database instance. Several criteria introducing sufficient conditions for chase termination have been recently proposed in the literature. The aim of this paper is to present more general criteria and techniques for chase termination. We first present extensions of the well-know stratification criterium and introduce a new criterium, called local stratification (LS), which generalizes both super-weak acyclicity and stratification-based criteria (including the class of constraints which are inductively restricted). Next the paper presents a rewriting algorithm, whose structure is similar to the one presented by Greco et al (2010); the algorithm takes as input a set of tuple generating dependencies and produces as output an equivalent set of dependencies and a boolean value stating whether a sort of cyclicity has been detected. The output set, obtained by adorning the input set of constraints, allows us to perform a more accurate analysis of the structural properties of constraints and to further enlarge the class of tuple generating dependencies is guaranteed, whereas the checking of acyclicity allows us to introduce the class of acyclic constraints (AC), which generalizes LS and guarantees chase termination.

Completeness of Queries over Incomplete Databases Simon Razniewski (FU Bozen), Werner Nutt (FU Bozen)

Data completeness is an important aspect of data quality as in many scenarios it is crucial to guarantee completeness of query answers. We develop techniques to conclude the completeness of query answers from information about the completeness of parts of a generally incomplete database. In our framework, completeness of a database can be described in two ways: by table completeness (TC) statements, which say that certain parts of a relation are complete, and by query completeness (QC) statements, which say that the set of answers of a query is complete. We identify as core problem to decide whether table completeness entails query completeness (TC-QC). We develop decision procedures and assess the complexity of TC-QC inferences depending on the languages of the TC and QC statements. We show that in important cases weakest preconditions for query completeness can be expressed in terms of table completeness statements, which means that these statements identify precisely the parts of a database that are critical for the completeness of a query. For the related problem of QC-QC entailment, we discuss its connection to query determinacy. Moreover, we show how to use the concrete state of a database to enable further completeness inferences.

Research 11: Graph Data

gStore: Answering SPARQL Queries via Subgraph Matching

Lei Zou (Peking University), Jinhui Mo (Peking University), Lei Chen (HKUST), M. Tamer Özsu (University of Waterloo), Zhao Dongyan (Peking University)

Due to the increasing use of RDF data, efficient processing of SPARQL queries over large RDF datasets has become an important issue. However, existing solutions suffer from two limitations: 1) they cannot answer SPARQL queries with wildcards in a scalable manner; and 2) they cannot handle frequent updates in RDF repositories efficiently. Thus, most of them have to reprocess the dataset from scratch. In this paper, we propose a graph-based approach to store and query RDF data. Rather than mapping RDF triples into a relational database as most existing methods do, we store RDF data as a large graph. A SPARQL query is then converted into a corresponding subgraph matching query. In order to speed up query processing, we develop a novel index, VS*-tree, together with some effective pruning rules and efficient searching algorithms. Our method can answer exact SPARQL queries and queries with wildcards in a uniform manner. Furthermore, we propose an effective maintenance algorithm to handle online updates over RDF repositories. Extensive experiments confirm the efficiency and effectiveness of our solution.

Ye Yuan (Northeastern University, China), Guoren Wang (Northeastern University, China), Haixun Wang (Microsoft Research Asia), Lei Chen (HKUST)

Retrieving graphs containing a query graph from a large graph database is a key task in many graph-based applications, including chemical compounds discovery, protein complex prediction, and structural pattern recognition. However, graph data handled by these applications is often noisy, incomplete, and inaccurate because of the way the data is produced. In this paper,we study subgraph queries over uncertain graphs. Specifically, we consider the problem of answering threshold-based probabilistic queries over a large uncertain graph database with the possible world semantics. We prove that problem is #P-complete, therefore, we adopt a filtering-and-verification strategy to speed up the search. In the filtering phase, we use a probabilistic inverted index, PIndex, based on subgraph features obtained by an optimal feature selection process. During the verification phase, we develop exact and bound algorithms to validate the remaining candidates. Extensive experimental results demonstrate the effectiveness of the proposed algorithms.

Scalable SPARQL Querying of Large RDF Graphs

Jiewen Huang (Yale University), Daniel Abadi (Yale University), Kun Ren (Northwestern Polytechnical University, China)

The generation of RDF data has accelerated to the point where many data sets need to be partitioned across multiple machines in order to achieve reasonable performance when querying the data. Although tremendous progress has been made in the Semantic Web community for achieving high performance data management on a single node, current solutions that allow the data to be partitioned across multiple machines are highly inefficient. In this paper, we introduce a scalable RDF data management system that is up to three orders of magnitude more efficient than popular multi-node RDF data management systems. In so doing, we introduce techniques for (1) leveraging state-of-the-art single node RDF-store technology (2) partitioning the data across nodes in a manner that helps accelerate query processing through locality optimizations and (3) decomposing SPARQL queries into high performance fragments that take advantage of how data is partitioned in a cluster.

Research 12: Cloud Computing and High-Availability

Albatross: Lightweight Elasticity in Shared Storage Databases for the Cloud using Live Data Migration

Sudipto Das (UC Santa Barbara), Shoji Nishimura (NEC Corporation), Divyakant Agrawal (UC Santa Barbara), Amr El Abbadi (UC Santa Barbara)

Database systems serving cloud platforms must serve large numbers of applications (or tenants). In addition to managing tenants with small data footprints, different schemas, and variable load patterns, such multitenant data platforms must minimize their operating costs by efficient resource sharing. When deployed over a pay-per-use infrastructure, elastic scaling and load balancing, enabled by low cost live migration of tenant databases, is critical to tolerate load variations while minimizing operating cost. However, existing databases - relational databases and Key-Value stores alike - lack low cost live migration techniques, thus resulting in heavy performance impact during elastic scaling. We present Albatross, a technique for live migration in a multitenant database serving OLTP style workloads where the persistent database image is stored in a network attached storage. Albatross migrates the database cache and the state of active transactions to ensure minimal impact on transaction execution while allowing transactions active during migration to continue execution. It also guarantees serializability while ensuring correctness during failures. Our evaluation using two OLTP benchmarks shows that Albatross can migrate a live tenant database with no aborted transactions, negligible impact on transaction latency and throughput both during and after migration, and an unavailability window as low as 300 ms.

iCBS: Incremental Cost-based Scheduling under Piecewise Linear SLAs

Yun Chi (NEC Labs America), Hyun Moon (NEC Labs America), Hakan Hacigumus (NEC Labs America)

In a cloud computing environment, it is beneficial for the cloud service provider to offer differentiated services among different customers, who often have different cost profiles. Therefore, cost-aware scheduling of queries is important. A practical cost-aware scheduling algorithm must be able to handle the highly demanding query volumes in the scheduling queues to make online scheduling decisions very quickly. We develop such a highly efficient cost-aware query scheduling algorithm, called iCBS. iCBS takes the query costs derived from the service level agreements (SLAs) between the service provider and its customers into account to make costaware scheduling decisions. iCBS is an incremental variation of an existing scheduling algorithm, CBS. Although CBS exhibits an exceptionally good cost performance, it has a prohibitive time complexity. Our main contributions are (1) to observe how CBS behaves under piecewise linear SLAs, which are very common in cloud computing systems, and (2) to efficiently leverage these observations and to reduce the online time complexity from O(N) for the original version CBS to O(log² N) for iCBS.

RemusDB: Transparent High Availability for Database Systems (Best Paper) Umar Farooq Minhas (University of Waterloo), Shriram Rajagopalan (University of

British Columbia), Brendan Cully (University of British Columbia), Ashraf Aboulnaga (University of Waterloo), Ken Salem (University of Waterloo), Andrew Warfield (University of British Columbia)

In this paper we present a technique for building a high-availability (HA) database management system (DBMS). The proposed technique can be applied to any DBMS with little or no customization, and with reasonable performance overhead. Our approach is based on Remus, a commodity HA solution implemented in the virtualization layer, that uses asynchronous virtual machine (VM) state replication to provide transparent HA and failover capabilities. We show that while Remus and similar systems can protect a DBMS, database workloads incur a performance overhead of up to 32% as compared to an unprotected DBMS. We identify the sources of this overhead and develop optimizations that mitigate the problems. We present an experimental evaluation using two popular database systems and industry standard benchmarks showing that for certain workloads, our optimized approach provides very fast failover (<= 3 seconds of downtime) with low performance overhead when compared to an unprotected DBMS. Our approach provides a practical means for existing, deployed database systems to be made more reliable with a minimum of risk, cost, and effort. Furthermore, this paper invites new discussion about whether the complexity of HA is best implemented within the DBMS, or as a service by the infrastructure below it.

Research 13: Human-Computer Interaction

SnipSuggest: Context-Aware Autocompletion for SQL

Nodira Khoussainova (University of Washington), YongChul Kwon (University of Washington), Magdalena Balazinska (University of Washington), Dan Suciu (University of Washington)

In this paper, we present SnipSuggest, a system that provides on-the-go, context-aware assistance in the SQL composition process. SnipSuggest aims to help the increasing population of non-expert database users, who need to perform complex analysis on their large-scale datasets, but have difficulty writing SQL. As a user types a query, it recommends possible additions

to various clauses in the query using relevant snippets collected from a log of past queries. SnipSuggest's current capabilities include suggesting tables, views, and table-valued functions in the FROM clause, columns in the SELECT clause, predicates in the WHERE clause, columns in the GROUP BY clause, aggregates, and some support for subqueries. SnipSuggest adjusts its recommendations according to the context: as the user writes more of the query, it is able to provide more accurate suggestions. We evaluate SnipSuggest over two query logs: one from an undergraduate database course and another from the SDSS database. We show that SnipSuggest is able to recommend useful snippets with up to 93.7% accuracy (average precision), at an interactive response time. We also show that SnipSuggest outperforms naive approaches, such as recommending the most popular snippets first.

A Probabilistic Approach for Automatically Filling Form-Based Web Interfaces Guilherme Toda (Federal University of Amazonas), Eli Cortez (Federal University of Amazonas), Altigran da Silva (Federal University of Amazonas), Edleno de Moura (Federal University of Amazonas)

In this paper we present a proposal for the implementation and evaluation of a novel method for automatically using data-rich text for filling form-based input interfaces. Our solution takes a text as input, extracts implicit data values from it and fills appropriate fields. For this task, we rely on knowledge obtained from values of previous submissions for each field, which are freely obtained from the usage of the interfaces. Our approach, called iForm, exploits features related to the content and the style of these values, which are combined through a Bayesian framework. Through extensive experimentation, we show that our approach is feasible and effective, and that it works well even when only a few previous submissions to the input interface are available.

Query Expansion Based on Clustered Results

Ziyang Liu (Arizona State University), Sivaramakrishnan Natarajan (Arizona State University), Yi Chen (Arizona State University)

Query expansion is a functionality of search engines that suggests a set of related queries for a user-issued keyword query. Typical corpus-driven keyword query expansion approaches return popular words in the results as expanded queries. Using these approaches, the expanded queries may correspond to a subset of possible query semantics, and thus miss relevant results. To handle ambiguous queries and exploratory queries, whose result relevance is difficult to judge, we propose a new framework for keyword query expansion: we start with clustering the results according to user specified granularity, and then generate expanded queries, such that one expanded query is generated for each cluster whose result set should ideally be the corresponding cluster. We formalize this problem and show its APX-hardness. Then we propose two efficient algorithms named iterative single-keyword refinement and partial elimination based convergence, respectively, which effectively generate a set of expanded queries from clustered results that provide a classification of the original query results. We believe our study of generating an optimal query based on the ground truth of the query results not only has applications in query expansion, but has significance for studying keyword search quality in general.

Research 14: Integrity Maintenance

Update Rewriting and Integrity Constraint Maintenance in a Schema Evolution Support System: PRISM++

Carlo Curino (MIT), Hyun Moon (NEC Labs America), Alin Deutsch (UCSD), Carlo Zaniolo (UCLA)

Supporting legacy applications when the database schema evolves represents a long-standing challenge of practical and theoretical importance. Recent work has produced algorithms and systems that automate the process of data migration and query adaptation; however, the

problems of evolving integrity constraints and supporting legacy updates under schema and integrity constraints evolution are significantly more difficult and have thus far remained unsolved. In this paper, we address this issue by introducing a formal evolution model for the database schema structure and its integrity constraints, and use it to derive update mapping techniques akin to the rewriting techniques used for queries. Thus, we (i) propose a new set of Integrity Constraints Modification Operators (ICMOs), (ii) characterize the impact on integrity constraints of structural schema changes, (iii) devise representations that enable the rewriting of updates, and (iv) develop a unified approach for query and update rewriting under constraints. We then describe the implementation of these techniques provided by our PRISM++ system. The effectiveness of PRISM++ and its enabling technology has been verified on a testbed containing evolution histories of several scientific databases and web information systems, including the Genetic DB Ensembl (410+ schema versions in 9 years), and Wikipedia (240+ schema versions in 6 years).

Business Policy Modeling and Enforcement in Databases Ahmed Ataullah (University of Waterloo), Frank Tompa (University of Waterloo)

Database systems are the central information repositories for businesses and are subject to a wide array of policies, rules and requirements. The spectrum of business level constraints implemented within database systems has expanded from classical access control to include auditing, usage control, privacy management, and records retention. The lack of a systematic mechanism of integrating and reasoning about such a diverse set of policies manifested as database level constraints makes corporate policy management a chaotic process. In this paper we propose a general purpose policy modeling and constraint management framework that can integrate numerous aspects of business level requirements within database systems. Our proposed solution relies on a finite state modeling language for business level policies, in which users can declaratively express rules related to the normal workflow of a business process as well as specifying any undesirable states of business objects contained in a database system. The proposed system is then able to translate these policies into low level temporal integrity constraints that prevent policy violations and ensure that business objects and artifacts follow their mandated lifecycles. A formal layer for reasoning allows policy makers to discover unenforceable and conflicting policies, providing the basis to guarantee compliance for a wide array of rules that may need to be enforced on complex business objects stored in relational database systems.

Data Coordination: Supporting Contingent Updates

Michael Lawrence (University of British Columbia), Rachel Pottinger (University of British Columbia), Sheryl Staub-French (University of British Columbia)

In many applications, a contingent data source may benefit by coordinating with external heterogeneous sources upon which it depends. The contingent source must be updated when changes are made do the external base sources; e.g. when a building design is updated, the contractor's cost estimate must be updated, too. The goal of data coordination is to update a contingent source, C, based on changes to an independently maintained base source B. This paper introduces a data coordination system allowing C to coordinate its data without imposing significant requirements on B. We use declarative mappings between B and C, and coordinate data in two stages: 1. View Differencing, which finds changes to an intermediate view of B based on its mapping to C, and 2. Update Translation, which translates the view differencing result into updates on C. We present and evaluate novel solutions to both stages and demonstrate their feasibility on real world problems.

Research 15: Distributed Systems

Using Paxos to Build a Scalable, Consistent, and Highly Available Datastore

Jun Rao (LinkedIn), Eugene Shekita (IBM Research - Almaden), Sandeep Tata (IBM Research - Almaden)

Spinnaker is an experimental datastore that is designed to run on a large cluster of commodity servers in a single datacenter. It features key-based range partitioning, 3-way replication, and a transactional get-put API with the option to choose either strong or timeline consistency on reads. This paper describes Spinnaker's Paxos-based replication protocol. The use of Paxos ensures that a data partition in Spinnaker will be available for reads and writes as long a majority of its replicas are alive. Unlike traditional master-slave replication, this is true regardless of the failure sequence that occurs. We show that Paxos replication can be competitive with alternatives that provide weaker consistency guarantees. Compared to an eventually consistent datastore, we show that Spinnaker can be as fast or even faster on reads and only 5% to 10% slower on writes.

A Framework for Supporting DBMS-like Indexes in the Cloud

Gang Chen (Zhejiang University), Hoang Tam Vo (National University of Singapore), Sai Wu (National University of Singapore), Beng Chin Ooi (National University of Singapore), M. Tamer Özsu (University of Waterloo)

To support "Database as a service" (DaaS) in the cloud, the database system is expected to provide similar functionalities as in centralized DBMS such as efficient processing of ad hoc queries. The system must therefore support DBMS-like indexes, possibly a few indexes for each table to provide fast location of data distributed over the network. In such a distributed environment, the indexes have to be distributed over the network to achieve scalability and reliability. Each cluster node maintains a subset of the index data. As in conventional DBMS, indexes incur maintenance overhead and the problem is more complex in the distributed environment since the data are typically partitioned and distributed based on a subset of attributes. Further, the distribution of indexes is not straight forward, and there is therefore always a question on scalability, in terms of data volume, network size, and number of indexes. In this paper, we examine the problem of providing DBMS-like indexing mechanisms in the cloud database systems, and propose an extensible, but simple and efficient indexing framework that enables users to define their own type of indexes without knowing the structure of the underlying network. It is also designed to ensure the efficiency of hopping between cluster nodes during index traversal, and reduce the maintenance cost of indexes. We implement three common indexes, namely distributed hash indexes, distributed B+-tree-like indexes and distributed multi-dimensional indexes, to demonstrate the usability and effectiveness of the framework. We conduct experiments on Amazon EC2 and an in-house cluster to verify the efficiency and scalability of the framework.

Serializable Snapshot Isolation for Replicated Databases in High-Update Scenarios Hyungsoo Jung (University of Sydney), Hyuck Han (Seoul National University), Alan Fekete (University of Sydney), Uwe Röhm (University of Sydney)

Many proposals for managing replicated data use sites running the Snapshot Isolation (SI) concurrency control mechanism, and provide 1-copy SI or something similar, as the global isolation level. This allows good scalability, since only ww-conflicts need to be managed globally. However, 1-copy SI can lead to data corruption and violation of integrity constraints. 1-copy serializability is the global correctness condition that prevents data corruption. We propose a new algorithm Replicated Serializable Snapshot Isolation (RSSI) that uses SI at each site, and combines this with a certification algorithm to guarantee 1-copy serializable global execution. Management of ww-conflicts is similar to what is done in 1-copy SI. But unlike previ-

ous designs for 1-copy serializable systems, we do not need to prevent all rw-conflicts among concurrent transactions. We formalize this in a theorem that shows that many rw-conflicts are indeed false positives that do not risk non-serializable behavior. Our proposed RSSI algorithm will only abort a transaction when it detects a well-defined pattern of two consecutive rw-edges in the serialization graph. We have built a prototype that integrates our RSSI with the existing open-source Postgres-R(SI) system. Our performance evaluation shows that there is a worst-case overhead of about 15% for getting full 1-copy serializability as compared to 1-copy SI in a cluster of 8 nodes, with our proposed RSSI clearly outperforming the previous work for update-intensive workloads.

Research 16: Streams and Events

Lightweight Graphical Models for Selectivity Estimation Without Independence Assumptions

Kostas Tzoumas (Aalborg University), Amol Deshpande (University of Maryland), Christian Jensen (Aarhus University)

As a result of decades of research and industrial development, modern query optimizers are complex software artifacts. However, the quality of the query plan chosen by an optimizer is largely determined by the quality of the underlying statistical summaries. Small selectivity estimation errors, propagated exponentially, can lead to severely sub-optimal plans. Modern optimizers typically maintain one-dimensional statistical summaries and make the attribute value independence and join uniformity assumptions for efficiently estimating selectivities. Therefore, selectivity estimation errors in today's optimizers are frequently caused by missed correlations between attributes. We present a selectivity estimation approach that does not make the independence assumptions. By carefully using concepts from the field of graphical models, we are able to factor the joint probability distribution of all the attributes in the database into small, usually two-dimensional distributions. We describe several optimizations that can make selectivity estimation highly efficient, and we present a complete implementation inside PostgreSQL's query optimizer. Experimental results indicate an order of magnitude better selectivity estimates, while keeping optimization time in the range of tens of milliseconds.

Active Complex Event Processing over Event Streams

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

State-of-the-art Complex Event Processing technology (CEP), while effective for pattern query execution, is limited in its capability of reacting to opportunities and risks detected by pattern queries. Especially reactions that affect the query results in turn have not been addressed in the literature. We propose to tackle these unsolved problems by embedding active rule support within the CEP engine, henceforth called Active CEP (ACEP). Active rules in ACEP allow us to specify a pattern query's dynamic condition and real-time actions. The technical challenge is to handle interactions between queries and reactions to queries in the high-volume stream execution. We hence introduce a novel stream-oriented transactional model along with a family of stream transaction scheduling algorithms that ensure the correctness of concurrent stream execution. We demonstrate the power of ACEP technology by applying it to the development of a healthcare system being deployed in UMass Medical School hospital. Through extensive performance experiments using real data streams, we show that our unique Low-Water-Mark stream transaction scheduler, customized for streaming environments, successfully achieves near-real-time system responsiveness and gives orders-of-magnitude better throughput than our alternative schedulers.

Massive Scale-out of Expensive Continuous Queries Erik Zeitler (Uppsala University), Tore Risch (Uppsala University)

Scalable execution of expensive continuous queries over massive data streams requires input streams to be split into parallel sub-streams. The query operators are continuously executed in parallel over these sub-streams. Stream splitting involves both partitioning and replication of incoming tuples, depending on how the continuous query is parallelized. We provide a stream splitting operator that enables such customized stream splitting. However, it is critical that the stream splitting predicates have some costs. Therefore, to enable customized splitting of high-volume streams, we introduce a parallelized stream splitting operator, called parasplit. We investigate the performance of parasplit using a cost model and experimentally. Based on these results, a heuristic is devised to automatically parallelize the execution of parasplit. We show that the maximum stream rate of parasplit is network bound, and that the parallelization is energy efficient. Finally, the scalability of our approach is experimentally demonstrated on the Linear Road Benchmark, showing an order of magnitude higher stream processing rate over previously published results, allowing at least 512 expressways.

Research 17: Privacy and Protection

Personalized Privacy Protection in Social Networks

Mingxuan Yuan (HKUST), Lei Chen (HKUST), Philip Yu (UIC)

Due to the popularity of social networks, many proposals have been proposed to protect the privacy of the networks. All these works assume that the attacks use the same background knowledge. However, in practice, different users have different privacy protect requirements. Thus, assuming the attacks with the same background knowledge does not meet the personalized privacy requirements, meanwhile, it looses the chance to achieve better utility by taking advantage of differences of users' privacy requirements. In this paper, we introduce a framework which provides privacy preserving services based on user's personal privacy requests. We define three levels of protection requirements based on the gradually increasing attacker's background knowledge and combine the label generalization protection and the structure protection techniques (i.e. adding noise edges or nodes) together to satisfy different users' protection requirements. We verify the effectiveness of the proposed framework through extensive experiments.

Publishing Set-Valued Data via Differential Privacy

Rui Chen (Concordia University), Noman Mohammed (Concordia University), Benjamin C. M. Fung (Concordia University), Bipin Desai (Concordia University), Li Xiong (Emory University)

Set-valued data provides enormous opportunities for various data mining tasks. In this paper, we study the problem of publishing set-valued data for data mining tasks under the rigorous differential privacy model. All existing data publishing methods for set-valued data are based on partition-based privacy models, for example k-anonymity, which are vulnerable to privacy attacks based on background knowledge. In contrast, differential privacy provides strong privacy guarantees independent of an adversary's background knowledge and computational power. Existing data publishing approaches for differential privacy, however, are not adequate in terms of both utility and scalability in the context of set-valued data due to its high dimensionality. We demonstrate that set-valued data could be efficiently released under differential privacy with guaranteed utility with the help of context-free taxonomy trees. We propose a probabilistic top-down partitioning algorithm to generate a differentially private release, which scales linearly with the input data size. We also discuss the applicability of our idea to the context of relational data. We prove that our result is (ε, δ) -useful for the class of counting queries,

the foundation of many data mining tasks. We show that our approach maintains high utility for counting queries and frequent itemset mining and scales to large datasets through extensive experiments on real-life set-valued datasets.

Private Analysis of Graph Structure

Vishesh Karwa (Pennsylvania State University), Sofya Raskhodnikova (Pennsylvania State University), Adam Smith (Pennsylvania State University), Grigory Yaroslavtsev (Pennsylvania State University)

We present efficient algorithms for releasing useful statistics about graph data while providing rigorous privacy guarantees. Our algorithms work on data sets that consist of relationships between individuals, such as social ties or email communication. The algorithms satisfy edge differential privacy, which essentially requires that the presence or absence of any particular relationship be hidden. Our algorithms output approximate answers to subgraph counting queries. Given a query graph H, e.g., a triangle, k-star or k-triangle, the goal is to return the number of edge-induced isomorphic copies of H in the input graph. The special case of triangles was considered by Nissim, Raskhodnikova and Smith (STOC 2007), and a more general investigation of arbitrary query graphs was initiated by Rastogi, Hay, Miklau and Suciu (PODS 2009). We extend the approach of [NRS] to a new class of statistics, namely, k-star queries. We also give algorithms for k-triangle queries using a different approach, based on the higher-order local sensitivity. For the specific graph statistics we consider (i.e., k-stars and k-triangles), we significantly improve on the work of [RHMS]: our algorithms satisfy a stronger notion of privacy, which does not rely on the adversary having a particular prior distribution on the data, and add less noise to the answers before releasing them. We evaluate the accuracy of our algorithms both theoretically and empirically, using a variety of real and synthetic data sets. We give explicit, simple conditions under which these algorithms add a small amount of noise. We also provide the average-case analysis in the Erdos-Renyi-Gilbert G(n,p) random graph model. Finally, we give hardness results indicating that the approach NRS used for triangles cannot easily be extended to k-triangles (and hence justifying our development of a new algorithmic approach).

Surrogate Parenthood: Protected and Informative Graphs Barbara Blaustein (MITRE), Adriane Chapman (MITRE), Len Seligman (MITRE), M. David Allen (MITRE), Arnon Rosenthal (MITRE)

Many applications, including provenance and some analyses of social networks, require pathbased queries over graph-structured data. When these graphs contain sensitive information, paths may be broken, resulting in uninformative query results. This paper presents innovative techniques that give users more informative graph query results; the techniques leverage a common industry practice of providing what we call surrogates: alternate, less sensitive versions of nodes and edges releasable to a broader community. We describe techniques for interposing surrogate nodes and edges to protect sensitive graph components, while maximizing graph connectivity and giving users as much information as possible. In this work, we formalize the problem of creating a protected account G' of a graph G. We provide a utility measure to compare the informativeness of alternate protected accounts and an opacity measure for protected accounts, which indicates the likelihood that an attacker can recreate the topology of the original graph from the protected account. We provide an algorithm to create a maximally useful protected account of a sensitive graph, and show through evaluation with the PLUS prototype that using surrogates and protected accounts adds value for the user, with no significant impact on the time required to generate results for graph queries.

Research 18: MapReduce and Hadoop

Column-Oriented Storage Techniques for MapReduce

Avrilia Floratou (University of Wisconsin-Madison), Jignesh Patel (University of Wisconsin-Madison), Eugene Shekita (IBM Research - Almaden), Sandeep Tata (IBM Research - Almaden)

Users of MapReduce often run into performance problems when they scale up their workloads. Many of the problems they encounter can be overcome by applying techniques learned from over three decades of research on parallel DBMSs. However, translating these techniques to a MapReduce implementation such as Hadoop presents unique challenges that can lead to new design choices. This paper describes how column-oriented storage techniques can be incorporated in Hadoop in a way that preserves its popular programming APIs. We show that simply using binary storage formats in Hadoop can provide a 3x performance boost over the naive use of text files. We then introduce a column-oriented storage format that is compatible with the replication and scheduling constraints of Hadoop and show that it can speed up MapReduce jobs on real workloads by an order of magnitude. We also show that dealing with complex column types such as arrays, maps, and nested records, which are common in MapReduce jobs, can incur significant CPU overhead. Finally, we introduce a novel skip list column format and lazy record construction strategy that avoids deserializing unwanted records to provide an additional 1.5x performance boost. Experiments on a real intranet crawl are used to show that our column-oriented storage techniques can improve the performance of the map phase in Hadoop by as much as two orders of magnitude.

Automatic Optimization for MapReduce Programs

Eaman Jahani (University of Michigan), Michael Cafarella (University of Michigan), Christopher Ré (University of Wisconsin-Madison)

The MapReduce distributed programming framework has become popular, despite evidence that current implementations are inefficient, requiring far more hardware than a traditional relational databases to complete similar tasks. MapReduce jobs are amenable to many traditional database query optimizations (B+Trees for selections, column-store-style techniques for projections, etc), but existing systems do not apply them, substantially because free-form user code obscures the true data operation being performed. For example, a selection in SQL is easily detected, but a selection in a MapReduce program is embedded in Java code along with lots of other program logic. We could ask the programmer to provide explicit hints about the program's data semantics, but one of MapReduce's attractions is precisely that it does not ask the user for such information. This paper covers Manimal, which automatically analyzes MapReduce programs and applies appropriate data-aware optimizations, thereby requiring no additional help at all from the programmer. We show that Manimal successfully detects optimization opportunities across a range of data operations, and that it yields speedups of up to 1,121% on previously-written MapReduce programs.

CoHadoop: Flexible Data Placement and Its Exploitation in Hadoop

Mohamed Eltabakh (IBM Research - Almaden), Yuanyuan Tian (IBM Research -Almaden), Fatma Ozcan (IBM Research - Almaden), Rainer Gemulla (Max-Planck-Institut für Informatik), Aljoscha Krettek (IBM Germany), John McPherson (IBM Research -Almaden)

Hadoop has become an attractive platform for large-scale data analytics. In this paper, we identify a major performance bottleneck of Hadoop: its lack of ability to colocate related data on the same set of nodes. To overcome this bottleneck, we introduce CoHadoop, a lightweight extension of Hadoop that allows applications to control where data are stored. In contrast to previous approaches, CoHadoop retains the flexibility of Hadoop in that it does not require us-

ers to convert their data to a certain format (e.g., a relational database or a specific file format). Instead, applications give hints to CoHadoop that some set of files are related and may be processed jointly; CoHadoop then tries to colocate these files for improved efficiency. Our approach is designed such that the strong fault tolerance properties of Hadoop are retained. Colocation can be used to improve the efficiency of many operations, including indexing, grouping, aggregation, columnar storage, joins, and sessionization. We conducted a detailed study of joins and sessionization in the context of log processing - a common use case for Hadoop -, and propose efficient map-only algorithms that exploit colocated data partitions. In our experiments, we observed that CoHadoop outperforms both plain Hadoop and previous work. Our approach not only performs better than repartition-based algorithms, but also outperforms map-only algorithms that do exploit data partitioning but not colocation.

Profiling, What-if Analysis, and Cost-based Optimization of MapReduce Programs Herodotos Herodotou (Duke University), Shivnath Babu (Duke University)

MapReduce has emerged as a viable competitor to database systems in big data analytics. MapReduce programs are being written for a wide variety of application domains including business data processing, text analysis, natural language processing, Web graph and social network analysis, and computational science. However, MapReduce systems lack a feature that has been key to the historical success of database systems, namely, cost-based optimization. A major challenge here is that, to the MapReduce system, a program consists of black-box map and reduce functions written in some programming language like C++, Java, Python, or Ruby. We introduce, to our knowledge, the first Cost-based Optimizer for simple to arbitrarily complex MapReduce programs. We focus on the optimization opportunities presented by the large space of configuration parameters for these programs. We also introduce a Profiler to collect detailed statistical information from unmodified MapReduce programs, and a What-if Engine for fine-grained cost estimation. All components have been prototyped for the popular Hadoop MapReduce system. The effectiveness of each component is demonstrated through a comprehensive evaluation using representative MapReduce programs from various application domains.

Research 19: Ranking

On Pruning for Top-K Ranking in Uncertain Databases

Chonghai Wang (University of Alberta), Li Yan Yuan (University of Alberta), Jia-Huai You (University of Alberta), Osmar Zaiane (University of Alberta), Jian Pei (Simon Fraser University)

Top-k ranking for an uncertain database is to rank tuples in it so that the best k of them can be determined. The problem has been formalized under the unified approach based on parameterized ranking functions (PRFs) and the possible world semantics. Given a PRF, one can always compute the ranking function values of all the tuples to determine the top-k tuples, which is a formidable task for large databases. In this paper, we present a general approach to pruning for the framework based on PRFs. We show a mathematical manipulation of possible worlds which reveals key insights in the part of computation that may be pruned and how to achieve it in a systematic fashion. This leads to concrete pruning methods for a wide range of ranking functions. We show experimentally the effectiveness of our approach.

PathSim: Meta Path-Based Top-K Similarity Search in Heterogeneous Information Networks

Yizhou Sun (UIUC), Jiawei Han (UIUC), Xifeng Yan (UCSB), Philip Yu (UIC), Tianyi Wu (Microsoft)

Similarity search is a primitive operation in database and Web search engines. With the advent of large-scale heterogeneous information networks that consist of multi-typed, interconnected

objects, such as the bibliographic networks and social media networks, it is important to study similarity search in such networks. Intuitively, two objects are similar if they are linked by many paths in the network. However, most existing similarity measures are defined for homogeneous networks. Different semantic meanings behind paths are not taken into consideration. Thus they cannot be directly applied to heterogeneous networks. In this paper, we study similarity search that is defined among the same type of objects in heterogeneous networks. Moreover, by considering different linkage paths in a network, one could derive various similarity semantics. Therefore, we introduce the concept of meta path-based similarity, where a meta path is a path consisting of a sequence of relations defined between different object types (i.e., structural paths at the meta level). No matter whether a user would like to explicitly specify a path combination given sufficient domain knowledge, or choose the best path by experimental trials, or simply provide training examples to learn it, meta path forms a common base for a network-based similarity search engine. In particular, under the meta path framework we define a novel similarity measure called PathSim that is able to find peer objects in the network (e.g., find authors in the similar field and with similar reputation), which turns out to be more meaningful in many scenarios compared with random-walk based similarity measures. In order to support fast online query processing for PathSim queries, we develop an efficient solution that partially materializes short meta paths and then concatenates them online to compute top-k results. Experiments on real data sets demonstrate the effectiveness and efficiency of our proposed paradigm.

Optimizing and Parallelizing Ranked Enumeration

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

Lawler-Murty's procedure is a general tool for designing algorithms for enumeration problems (i.e., problems that involve the production of a large set of answers in ranked order), which naturally arise in database management. Lawler-Murty's procedure is used in a variety of modern database applications; particularly in those related to keyword search over structured data. Essentially, this procedure enumerates by invoking a series of instances of an optimization problem (i.e., finding the best solution); solving the optimization problem is the only part that depends on the specific task at hand. The topic of optimizing and parallelizing Lawler-Murty's procedure is investigated. Naive parallelism can be carried out by concurrently solving independent instances of the optimization problem. This can be improved by printing the next answer, in the enumeration order, as soon as none of the concurrent instances can produce a better answer. However, this approach alone suffers from poor utilization of available threads. That leads to the idea of freezing an instance of the optimization problem. Interestingly, not only is freezing beneficial to the parallel execution of Lawler-Murty's, it also substantially reduces the running time of the serial execution. Additional improvements of the freezing technique are then developed to further enhance the utilization of threads, and they result in a significant overall speedup. The effectiveness of the proposed approach is demonstrated on keyword search over data graphs, wherein an extensive experimental study is described.

Efficient Rank Join with Aggregation Constraints

Min Xie (University of British Columbia), Laks Lakshmanan (University of British Columbia), Peter Wood (Birkbeck, University of London)

We show aggregation constraints that naturally arise in several applications can enrich the semantics of rank join queries, by allowing users to impose their application-specific preferences in a declarative way. By analyzing the properties of aggregation constraints, we develop efficient deterministic and probabilistic algorithms which can push the aggregation constraints inside the rank join framework. Through extensive experiments on various datasets, we show that in many cases our proposed algorithms can significantly outperform the naive approach of applying the state-of-the-art rank join algorithm followed by post-filtering to discard results violating the constraints.

Research 20: Statistical Methods

Tuffy: Scaling up Statistical Inference in Markov Logic Networks using an RDBMS Feng Niu (University of Wisconsin-Madison), Christopher Ré (University of Wisconsin-Madison), AnHai Doan (University of Wisconsin-Madison), Jude Shavlik (University of Wisconsin-Madison)

Markov Logic Networks (MLNs) have emerged as a powerful framework that combines statistical and logical reasoning; they have been applied to many data intensive problems including information extraction, entity resolution, and text mining. Current implementations of MLNs do not scale to large real-world data sets, which is preventing their wide-spread adoption. We present Tuffy that achieves scalability via three novel contributions: (1) a bottom-up approach to grounding that allows us to leverage the full power of the relational optimizer, (2) a novel hybrid architecture that allows us to perform AI-style local search efficiently using an RDBMS, and (3) a theoretical insight that shows when one can (exponentially) improve the efficiency of stochastic local search. We leverage (3) to build novel partitioning, loading, and parallel algorithms. We show that our approach outperforms state-of-the-art implementations in both quality and speed on several publicly available datasets.

Dissemination of Models over Time-Varying Data

Yongluan Zhou (University of Southern Denmark), Zografoula Vagena (Rice University), Jonas Haustad (University of Southern Denmark)

Dissemination of time-varying data is essential in many applications, such as sensor networks, patient monitoring, stock tickers, etc. Often, the raw data have to go through some form of pre-processing, such as cleaning, smoothing, etc, before being disseminated. Such pre-processing often applies mathematical or statistical models to transform the large volumes of raw, point-based data into a much smaller number of piece-wise continuous functions. In such cases, the necessity to distribute data models instead of raw data may arise. Nevertheless, model dissemination has received very little attention so far. In this paper, we attempt to fill this gap and propose a model-agnostic dissemination framework that can handle different models in a uniform manner. The dissemination infrastructure is built on top of a tree-based overlay network, reminiscent to the ones employed in publish/subscribe systems, which are known to scale well to the number of data producers and receivers. To adequately deal with the vast model variation and receivers' very different accuracy requirements on the models, we have developed optimized model routing algorithms, which are intended to minimize data traffic and avoid bottlenecks within the dissemination network. The extensive experimental evaluation over a prototype system that we have built shows that our methods are both effective and robust.

Storing Matrices on Disk: Theory and Practice Revisited

Yi Zhang (Duke University), Kamesh Munagala (Duke University), Jun Yang (Duke University)

We consider the problem of storing arrays on disk to support scalable data analysis involving linear algebra. We propose Linearized Array B-tree, or LAB-tree, which supports flexible array layouts and automatically adapts to varying sparsity across parts of an array and over time. We reexamine the B-tree splitting strategy for handling insertions and the flushing policy for batching updates, and show that common practices may in fact be suboptimal. Through theoretical and empirical studies, we propose alternatives with good theoretical guarantees and/or practical performance.

Online Aggregation for Large MapReduce Jobs

Niketan Pansare (Rice University), Vinayak Borkar (UC Irvine), Chris Jermaine (University of Florida), Tyson Condie (Yahoo! Research)

In online aggregation, a database system processes a user's aggregation query in an online fashion. At all times during processing, the system gives the user an estimate of the final query result, with the confidence bounds that become tighter over time. In this paper, we consider how online aggregation can be built into a MapReduce system for large-scale data processing. Given the MapReduce paradigm's close relationship with cloud computing (in that one might expect a large fraction of MapReduce jobs to be run in the cloud), online aggregation is a very attractive technology. Since large-scale cloud computations are typically pay-as-you-go, a user can monitor the accuracy obtained in an online fashion, and then save money by killing the computation early once sufficient accuracy has been obtained.

Research 21: Graph Data

Distance-Constraint Reachability Computation in Uncertain Graphs

Ruoming Jin (Kent State University), Lin Liu (Kent State University), Bolin Ding (UIUC), Haixun Wang (Microsoft Research Asia)

Driven by the emerging network applications, querying and mining uncertain graphs has become increasingly important. In this paper, we investigate a fundamental problem concerning uncertain graphs, which we call the distance-constraint reachability (DCR) problem: Given two vertices s and t, what is the probability that the distance from s to t is less than or equal to a user-defined threshold d in the uncertain graph? Since this problem is #P-Complete, we focus on efficiently and accurately approximating DCR online. Our main results include two new estimators for the probabilistic reachability. One is a Horvitz-Thomson type estimator based on the unequal probabilistic sampling scheme, and the other is a novel recursive sampling estimator, which effectively combines a deterministic recursive computational procedure with a sampling process to boost the estimation accuracy. Both estimators can produce much smaller variance than the direct sampling estimator, which considers each trial to be either 1 or 0. We also present methods to make these estimators more computationally efficient. The comprehensive experiment evaluation on both real and synthetic datasets demonstrates the efficiency and accuracy of our new estimators.

Keyword Search in Graphs: Finding r-cliques Mehdi Kargar (York University), Aijun An (York University)

Keyword search over a graph finds a substructure of the graph containing all or some of the input keywords. Most of previous methods in this area find connected minimal trees that cover all the query keywords. Recently, it has been shown that finding subgraphs rather than trees can be more useful and informative for the users. However, the current tree or graph based methods may produce answers in which some content nodes (i.e., nodes that contain input keywords) are not very close to each other. In addition, when searching for answers, these methods may explore the whole graph rather than only the content nodes. This may lead to poor performance in execution time. To address the above problems, we propose the problem of finding r-cliques in graphs. An r-clique is a group of content nodes that cover all the input keywords and the distance between each two nodes is less than or equal to r. An exact algorithm that produces r-cliques with 2-approximation in polynomial delay is proposed. Extensive performance studies using two large real data sets confirm the efficiency and accuracy of finding r-cliques in graphs.

On Link-based Similarity Join

Liwen Sun (University of Hong Kong), Reynold Cheng (University of Hong Kong), Xiang Li (University of Hong Kong), David Cheung (University of Hong Kong), Jiawei Han (UIUC)

Graphs can be found in applications like social networks, bibliographic networks, and biological databases. Understanding the relationship, or links, among graph nodes enables applications such as link prediction, recommendation, and spam detection. In this paper, we propose link-based similarity join (LS-join), which extends the similarity join operator to link-based measures. Given two sets of nodes in a graph, the LS-join returns all pairs of nodes that are highly similar to each other, with respect to an e-function. The e-function generalizes common measures like Personalized PageRank (PPR) and SimRank (SR). We study an efficient LS-join algorithm on a large graph. We further improve our solutions for PPR and SR, which involve expensive random-walk operations. We validate our solutions by performing extensive experiments on three real graph datasets.

Research 22: Data Integration

Synthesizing Products for Online Catalogs

Hoa Nguyen (University of Utah), Ariel Fuxman (Microsoft Research), Stelios Paparizos (Microsoft Research), Juliana Freire (University of Utah), Rakesh Agrawal (Microsoft Research)

A high-quality, comprehensive product catalog is essential to the success of Product Search engines and shopping sites such as Yahoo! Shopping, Google Product Search, and Bing Shopping. Given the large number of products and the speed at which they are released to the market, keeping catalogs up-to-date becomes a challenging task, calling for the need of automated techniques. In this paper, we introduce the problem of product synthesis, a key component of catalog creation and maintenance. Given a set of offers advertised by merchants, the goal is to identify new products and add them to the catalog, together with their (structured) attributes. A fundamental challenge in product synthesis is the scale of the problem: a Product Search engine receives data from thousands of merchants and millions of products; the product taxonomy contains thousands of categories, where each category has a different schema; and merchants use representations for products that are different from the ones used in the catalog of the Product Search engine. We propose a system that provides an end-to-end solution to the product synthesis problem, and includes components for extraction, and addresses issues involved in data extraction from offers, schema reconciliation, and data fusion. We developed a novel and scalable technique for schema matching which leverages knowledge about previously-known instance-level associations between offers and products; and it is trained using automatically created training sets (no manually-labeled data is needed). We present an experimental evaluation of our system using data from Bing Shopping for more than 800K offers, a thousand merchants, and 400 categories. The evaluation confirms that our approach is able to automatically generate a large number of accurate product specifications. Furthermore, the evaluation shows that our schema reconciliation component outperforms state-of-the-art schema matching techniques in terms of precision and recall.

Online Data Fusion

Xuan Liu (National University of Singapore), Xin Dong (AT&T Labs), Beng Chin Ooi (National University of Singapore), Divesh Srivastava (AT&T Labs)

The Web contains a significant volume of structured data in various domains, but meanwhile, a fair volume of Web data are dirty and erroneous, and such data can get quickly populated through copying between Web sources. While advanced data integration techniques allow querying structured data on the Web, they return the union of the answers retrieved from different sources and can thus return conflicting information, leaving decisions of which answers are

correct to the end users. Data fusion techniques, on the other hand, aim at finding the true values that reflect the real world, but are designed for offline data aggregation and can take a long time. This paper proposes Solaris, the first online data fusion system. It starts with returning answers from the first probed source, and refreshes the answers as it probes more sources and applies fusion techniques on the retrieved data. For each returned answer, it shows the likelihood that the answer is correct, and terminates after gaining enough confidence that data from the unprocessed sources are unlikely to change the returned answers. We address key problems in building such a system and show empirically that the system can return correct answers quickly and probe much less sources without sacrificing the quality of the answers.

Research 23: Social Networks

Structural Trend Analysis for Online Social Networks

Ceren Budak (UC Santa Barbara), Divyakant Agrawal (UC Santa Barbara), Amr El Abbadi (UC Santa Barbara)

The identification of popular and important topics discussed in social networks is crucial for a better understanding of societal concerns. It is also useful for users to stay on top of trends without having to sift through vast amounts of shared information. Trend detection methods introduced so far have not used the network topology and has thus not been able to distinguish viral topics from topics that are diffused mostly through the news media. To address this gap, we propose two novel structural trend definitions we call coordinated and uncoordinated trends that use friendship information to identify topics that are discussed among clustered and distributed users respectively. Our analyses and experiments show that structural trends are significantly different from traditional trends and provide new insights into the way people share information online. We also propose a sampling technique for structural trend detection and prove that the solution yields in a gain in efficiency and is within an acceptable error bound. Experiments performed on a Twitter data set of 41.7 million nodes and 417 million posts show that even with a sampling rate of 0.005, the average precision is 0.93 for coordinated trends and 1 for uncoordinated trends.

On Social-Temporal Group Query with Acquaintance Constraint

De-Nian Yang (Academia Sinica), Yi-Ling Chen (National Taiwan University), Wang-Chien Lee (Pennsylvania State University), Ming-Syan Chen (National Taiwan University)

Three essential criteria are important for activity planning, including: (1) finding a group of attendees familiar with the initiator, (2) ensuring each attendee in the group to have tight social relations with most of the members in the group, and (3) selecting an activity period available for all attendees. Therefore, this paper proposes Social-Temporal Group Query to find the activity time and attendees with the minimum total social distance to the initiator. Moreover, this query incorporates an acquaintance constraint to avoid finding a group with mutually unfamiliar attendees. Efficient processing of the social-temporal group query is very challenging. We show that the problem is NP-hard via a proof and formulate the problem with Integer Programming. We then propose two efficient algorithms, SGSelect and STGSelect, which include effective pruning techniques and employ the idea of pivot time slots to substantially reduce the running time, for finding the optimal solutions. Experimental results indicate that the proposed algorithms are much more efficient and scalable. In the comparison of solution quality, we show that STGSelect outperforms the algorithm that represents manual coordination by the initiator.

Social Content Matching in MapReduce

Gianmarco De Francisci Morales (IMT Lucca), Aristides Gionis (Yahoo! Research), Mauro Sozio (MPI)

Matching problems are ubiquitous. They occur in economic markets, labor markets, internet advertising, and elsewhere. In this paper we focus on an application of matching for social

VLDB 2011

media. Our goal is to distribute content from information suppliers to information consumers. We seek to maximize the overall relevance of the matched content from suppliers to consumers while regulating the overall activity, e.g., ensuring that no consumer is overwhelmed with data and that all suppliers have chances to deliver their content. We propose two matching algorithms, GreedyMR and StackMR, geared for the MapReduce paradigm. Both algorithms have provable approximation guarantees, and in practice they produce high-quality solutions. While both algorithms scale extremely well, we can show that StackMR requires only a polylogarithmic number of MapReduce steps, making it an attractive option for applications with very large datasets. We experimentally show the trade-offs between quality and efficiency of our solutions on two large datasets coming from real-world social-media web sites.

Research 24: Searching and Ranking

Fast Incremental and Personalized PageRank

Bahman Bahmani (Stanford University), Abdur Chowdhury (Twitter Inc.), Ashish Goel (Stanford University, Twitter Inc.)

In this paper, we analyze the efficiency of Monte Carlo methods for incremental computation of PageRank, personalized PageRank, and similar random walk based methods (with focus on SALSA), on large-scale dynamically evolving social networks. We assume that the graph of friendships is stored in distributed shared memory, as is the case for large social networks such as Twitter. For global PageRank, we assume that the social network has n nodes, and m adversarially chosen edges arrive in a random order. We show that with a reset probability of ε , the expected total work needed to maintain an accurate estimate (using the Monte Carlo method) of the PageRank of every node at all times is $O(n \ln(m) / \epsilon^2)$. This is significantly better than all known bounds for incremental PageRank. For instance, if we naively recompute the PageRanks as each edge arrives, the simple power iteration method needs $\Omega(m^2/\ln(1/(1-\epsilon)))$ total time and the Monte Carlo method needs $O(mn/\epsilon)$ total time; both are prohibitively expensive. Furthermore, we also show that we can handle deletions equally efficiently. We then study the computation of the top k personalized PageRanks starting from a seed node, assuming that personalized PageRanks follow a power-law with exponent $\alpha < 1$. We show that if we store R>q In n random walks starting from every node for large enough constant q (using the approach outlined for global PageRank), then the expected number of calls made to the distributed social network database is O($k / (R^{(1-\alpha)})$). We also present experimental results from the social networking site, Twitter, verifying our assumptions and analyses. The overall result is that this algorithm is fast enough for real-time queries over a dynamic social network.

Efficient Diversification of Web Search Results

Gabriele Capannini (ISTI-CNR), Franco Maria Nardini (ISTI-CNR), Raffaele Perego (ISTI-CNR), Fabrizio Silvestri (ISTI-CNR)

In this paper we analyze the efficiency of various search results diversification methods. While efficacy of diversification approaches has been deeply investigated in the past, response time and scalability issues have been rarely addressed. A unified framework for studying performance and feasibility of result diversification solutions is thus proposed. First we define a new methodology for detecting when, and how, query results need to be diversified. To this purpose, we rely on the concept of "query refinement" to estimate the probability of a query to be ambiguous. Then, relying on this novel ambiguity detection method, we deploy and compare on a standard test set, three different diversification methods: IASelect, xQuAD, and OptSelect. While the first two are recent state-of-the-art proposals, the latter is an original algorithm introduced in this paper. We evaluate both the efficiency and the effectiveness of our approach against its competitors by using the standard TREC Web diversification track testbed. Results shown that OptSelect is able to run two orders of magnitude faster than the two other state-of-the-art approaches and to obtain comparable figures in diversification effectiveness.

Keyword Search on Form Results

Aditya Ramesh (Stanford University), S. Sudarshan (IIT Bombay), Purva Joshi (IIT Bombay)

In recent years there has been a good deal of research in the area of keyword search on structured and semi-structured data. Most of this body of work has a significant limitation in the context of enterprise data since it ignores the application code that has often been carefully designed to present data in a meaningful fashion to users. In this work, we consider how to perform keyword search on enterprise applications, which provide a number of forms that can take parameters; parameters may be explicit, or implicit such as the identifier of the user. In the context of such applications, the goal of keyword search is, given a set of keywords, to retrieve forms, along with parameter values, such that result of each retrieved form executed on the corresponding retrieved parameter values will contain the specified keywords. Some earlier work in this area was based on creating keyword indices on form results, but there are problems in maintaining such indices in the face of updates. In contrast, we propose techniques based on creating inverted SQL queries from the SQL queries in the forms. Unlike earlier work, our techniques do not require any special purpose indices, and instead make use of standard text indices supported by most database systems. We have implemented our techniques and show that keyword search can run at reasonable speeds even on large databases with a significant number of forms.

Research 25: Statistical Methods

Incrementally Maintaining Classification using an RDBMS Mehmet Levent Koc (University of Wisconsin-Madison), Christopher Ré (University of Wisconsin-Madison)

The proliferation of imprecise data has motivated both researchers and the database industry to push statistical techniques into relational database management systems (RDBMSs). We study algorithms to maintain model-based views for a popular statistical technique, classification, inside an RDBMS in the presence of updates to the training examples. We make three technical contributions: (1) An algorithm that incrementally maintains classification inside an RD-BMS. (2) An analysis of the above algorithm that shows that our algorithm is optimal among all deterministic algorithms (and asymptotically within a factor of 2 of a nondeterministic optimal). (3) An index structure based on the technical ideas that underlie the above algorithm which allows us to store only a fraction of the entities in memory. We apply our techniques to text processing, and we demonstrate that our algorithms provide several orders of magnitude improvement over non-incremental approaches to classification on a variety of data sets: such as the Cora, UCI Machine Learning Repository data sets, Citeseer, and DBLife.

An Incremental Hausdorff Distance Calculation Algorithm

Sarana Nutanong (University of Maryland), Edwin Jacox (University of Maryland), Hanan Samet (University of Maryland)

The Hausdorff distance is commonly used as a similarity measure between two point sets. Using this measure, a set X is considered similar to Y iff every point in X is close to at least one point in Y. Formally, the Hausdorff distance HausDist(X,Y) can be computed as the Max-Min distance from X to Y, i.e., find the maximum of the distance from an element in X to its nearest neighbor (NN) in Y. Although this is similar to the closest pair and farthest pair problems, computing the Hausdorff distance is a more challenging problem since its Max-Min nature involves both maximization and also minimization rather than just one or the other. A traditional approach to computing HausDist(X,Y) performs a linear scan over X and utilizes an index to help compute the NN in Y for each x in X. We present a pair of basic solutions that avoid scanning X by applying the concept of aggregate NN search to searching for the element in X that yields the Hausdorff distance. In addition, we propose a novel method which incrementally explores the indexes of the two sets X and Y simultaneously. As an example application of our techniques, we use the Hausdorff distance as a measure of similarity between two trajectories (represented as point sets). We also use this example application to compare the performance of our proposed method with the traditional approach and the basic solutions. Experimental results show that our proposed method outperforms all competitors by one order of magnitude in terms of the tree traversal cost and total response time.

Structure-Aware Sampling: Flexible and Accurate Summarization Edith Cohen (AT&T Labs), Graham Cormode (AT&T Labs), Nick Duffield (AT&T Labs)

In processing large quantities of data, a fundamental problem is to obtain a summary which supports approximate query answering. Random sampling yields flexible summaries which naturally support subset-sum queries with unbiased estimators and well-understood confidence bounds. Classic sample-based summaries, however, are designed for arbitrary subset queries and are oblivious to the structure in the set of keys. The particular structure, such as hierarchy, order, or product space (multi-dimensional), make range queries much more relevant for most analysis of the data. Dedicated summarization algorithms for range-sum queries have also been extensively studied. They can outperform existing sampling schemes in terms of accuracy on range queries per summary size. Their accuracy, however, rapidly degrades when, as is often the case, the query spans multiple ranges. They are also less flexible - being targeted for range sum queries alone - and are often quite costly to build and use. In this paper we propose and evaluate variance optimal sampling schemes that are structure-aware. These summaries improve over the accuracy of existing structure-oblivious sampling schemes on range queries while retaining the benefits of sample-based summaries: flexible summaries, with high accuracy on both range queries and arbitrary subset queries.

Research 26: Recommender Systems

Personalized Social Recommendations - Accurate or Private?

Ashwin Machanavajjhala (Yahoo! Research), Aleksandra Korolova (Stanford University), Atish Das Sarma (Google Research)

With the recent surge of social networks like Facebook, new forms of recommendations have become possible - personalized recommendations of ads, content, and even new friend and product connections based on one's social interactions. Since recommendations may use sensitive social information, it is speculated that these recommendations are associated with privacy risks. The main contribution of this work is in formalizing these expected trade-offs between the accuracy and privacy of personalized social recommendations. In this paper, we study whether "social recommendations", or recommendations that are solely based on a user's social network, can be made without disclosing sensitive links in the social graph. More precisely, we quantify the loss in utility when existing recommendation algorithms are modified to satisfy a strong notion of privacy, called differential privacy. We prove lower bounds on the minimum loss in utility for any recommendation algorithm that is differentially private. We adapt two privacy preserving algorithms from the differential privacy literature to the problem of social recommendations, and analyze their performance in comparison to the lower bounds, both analytically and experimentally. We show that good private social recommendations are feasible only for a small subset of the users in the social network or for a lenient setting of privacy parameters.

RecBench: Benchmarks for Evaluating Performance of Recommender System Architectures

Justin Levandoski (University of Minnesota), Michael Ekstrand (University of Minnesota), Michael Ludwig (University of Minnesota), Ahmed Eldawy (University of Minnesota), Mohamed Mokbel (University of Minnesota), John Riedl (University of Minnesota)

Traditionally, recommender systems have been "hand-built", implemented as custom applications hard-wired to a particular recommendation task. Recently, the database community has begun exploring alternative DBMS-based recommender system architectures, whereby a database both stores the recommender system data (e.g., ratings data and the derived recommender models) and generates recommendations using SQL queries. In this paper, we present a comprehensive experimental comparison of both architectures. We define a set of benchmark tasks based on the needs of a typical recommender-powered e-commerce site. We then evaluate the performance of the "hand-built" MultiLens recommender application against two DBMSbased implementations: an unmodified DBMS and RecStore, a DBMS modified to improve efficiency in incremental recommender model updates. We employ two non-trivial data sets in our study: the 10 million rating MovieLens data, and the 100 million rating data set used in the Netflix Challenge. This study is the first of its kind, and our findings reveal an interesting trade-off: "hand-built" recommenders exhibit superior performance in model-building and pure recommendation tasks, while DBMS-based recommenders are superior at more complex recommendation tasks such as providing filtered recommendations and blending text-search with recommendation prediction scores.

MRI: Meaningful Interpretations of Collaborative Ratings

Mahashweta Das (University of Texas, Arlington), Sihem Amer-Yahia (Yahoo! Research), Gautam Das (University of Texas, Arlington), Cong Yu (Google Research)

Collaborative rating sites (e.g., IMDb and CNet) have become essential resources that many users consult to make purchasing decisions on various items such as movies and products. Ideally, a user wants to quickly decide whether an item is desirable, especially when many choices are available. In practice, however, a user either spends a lot of time examining ratings and reviews before making an informed decision, or simply trusts overall rating aggregations associated with an item. In this paper, we argue that neither option is satisfactory and propose a novel and powerful third option, Meaningful Ratings Interpretation (MRI), that automatically provides a meaningful interpretation of ratings associated with the input items. As a simple example, given the movie "Usual Suspects," instead of simply showing the average rating of 8.7 from all reviewers, MRI produces a set of meaningful factoids such as "male reviewers under 30 from NYC love this movie". We define the notion of meaningful interpretation based on the idea of data cube, and formalize two important sub-problems, meaningful description mining and meaningful difference mining. We show that these problems are NP-hard and design novel randomized hill exploration algorithms to solve them efficiently. We conduct user studies to show that MRI provides more helpful information to the users than simple average ratings. Performance evaluation over real data shows that our algorithms perform much faster and generate equally good interpretations as brute-force algorithms.

Research 27: GPU-based Architectures and Column-store Indexing

Fast Sparse Matrix-Vector Multiplication on GPUs: Implications for Graph Mining Xintian Yang (Ohio State University), Srinivasan Parthasarathy (Ohio State University), Ponnuswamy Sadayappan (Ohio State University)

Scaling up the sparse matrix-vector multiplication kernel on modern Graphics Processing Units (GPU) has been at the heart of numerous studies in both academia and industry. In this article we present a novel non-parametric, self-tunable, approach to data representation for computing this kernel, particularly targeting sparse matrices representing power-law graphs. Using real data, we show how our representation scheme, coupled with a novel tiling algorithm, can yield significant benefits over the current state of the art GPU efforts on a number of core data mining algorithms such as PageRank, HITS and Random Walk with Restart.

High-Throughput Transaction Executions on Graphics Processors

Bingsheng He (Nanyang Technological University), Jeffrey Xu Yu (Chinese University of Hong Kong)

OLTP (On-Line Transaction Processing) is an important business system sector in various traditional and emerging online services. Due to the increasing number of users, OLTP systems require high throughput for executing tens of thousands of transactions in a short time period. Encouraged by the recent success of GPGPU (General-Purpose computation on Graphics Processors), we propose GPUTx, an OLTP engine performing high-throughput transaction executions on the GPU for in-memory databases. Compared with existing GPGPU studies usually optimizing a single task, transaction executions require handling many small tasks concurrently. Specifically, we propose the bulk execution model to group multiple transactions into a bulk and to execute the bulk on the GPU as a single task. The transactions within the bulk are executed concurrently on the GPU. We study three basic execution strategies (one with locks and the other two lock-free), and optimize them with the GPU features including the hardware support of atomic operations, the massive thread parallelism and the SPMD (Single Program Multiple Data) execution. We evaluate GPUTx on a recent NVIDIA GPU in comparison with its counterpart on a quad-core CPU. Our experimental results show that optimizations on GPUTx significantly improve the throughput, and the optimized GPUTx achieves 4-10 times higher throughput than its CPU-based counterpart on public transaction processing benchmarks.

Efficient Parallel Lists Intersection and Index Compression Algorithms using Graphics Processing Units

Naiyong Ao (Nankai University), Fan Zhang (Nankai University), Di Wu (Nankai University), Douglas Stones (Monash University), Gang Wang (Nankai University), Xiaoguang Liu (Nankai University), Jing Liu (Nankai University), Sheng Lin (Nankai University)

Major web search engines answer thousands of queries per second requesting information about billions of web pages. The data sizes and query loads are growing at an exponential rate. To manage the heavy workload, we consider techniques for utilizing a Graphics Processing Unit (GPU). We investigate new approaches to improve two important operations of search engines - lists intersection and index compression. For lists intersection, we develop techniques for efficient implementation of the binary search algorithm for parallel computation. We inspect some representative real-world datasets and find that a sufficiently long inverted list has an overall linear rate of increase. Based on this observation, we propose Linear Regression and Hash Segmentation techniques for contracting the search range. For index compression, the traditional d-gap based compression schemata are not well-suited for parallel computation, so we propose a Linear Regression Compression schema which has an inherent parallel structure. We further discuss how to efficiently intersect the compressed lists on a GPU. Our experimental results show significant improvements in the query processing throughput on several datasets.

Merging What's Cracked, Cracking What's Merged: Adaptive Indexing in Main-Memory Column-Stores

Stratos Idreos (CWI), Stefan Manegold (CWI), Harumi Kuno (HP Labs), Goetz Graefe (HP Labs)

Adaptive indexing is characterized by the partial creation and refinement of the index as side effects of query execution. Dynamic or shifting workloads may benefit from preliminary index

structures focused on the columns and specific key ranges actually queried - without incurring the cost of full index construction. The costs and benefits of adaptive indexing techniques should therefore be compared in terms of initialization costs, the overhead imposed upon queries, and the rate at which the index converges to a state that is fully-refined for a particular workload component. Based on an examination of database cracking and adaptive merging, which are two techniques for adaptive indexing, we seek a hybrid technique that has a low initialization cost and also converges rapidly. We find the strengths and weaknesses of database cracking and adaptive merging complementary. One has a relatively high initialization cost but converges rapidly. The other has a low initialization cost but converges relatively slowly. We analyze the sources of their respective strengths and explore the space of hybrid techniques. We have designed and implemented a family of hybrid algorithms in the context of a full in-memory database system. Our experiments compare their behavior against database cracking and adaptive merging, as well as against both traditional full index lookup and scan of unordered data. We show that the new hybrids significantly improve over past methods while at least two of the hybrids come very close to the "ideal performance" in terms of both overhead per query and convergence to a final state.

Research 28: Transaction Processing

PLP: Page Latch-free Shared-everything OLTP

Ippokratis Pandis (Carnegie Mellon University), Pinar Tözün (EPFL), Ryan Johnson (University of Toronto), Anastasia Ailamaki (EPFL)

Scaling the performance of shared-everything transaction processing systems to highly-parallel multicore hardware remains a challenge for database system designers. Recent proposals alleviate locking and logging bottlenecks in the system, leaving page latching as the next potential problem. To tackle the page latching problem, we propose physiological partitioning (PLP). The PLP design applies logical-only partitioning, maintaining the desired properties of sharedeverything designs, and introduces a multi-rooted B+Tree index structure (MRBTree) which enables the partitioning of the accesses at the physical page level. Logical partitioning and MRBTrees together ensure that all accesses to a given index page come from a single thread and, hence, can be entirely latch-free; an extended design makes heap page accesses thread-private as well. Eliminating page latching allows us to simplify key code paths in the system such as B+Tree operations leading to more efficient and maintainable code. Profiling a prototype PLP system running on different multicore machines shows that it acquires 85% and 68% fewer contentious critical sections, respectively, than an optimized conventional design and one based on logical-only partitioning. PLP also improves performance up to 40% and 18%, respectively, over the existing systems.

Implementing Performance Competitive Logical Recovery

David Lomet (Microsoft Research), Kostas Tzoumas (Aalborg University), Michael Zwilling (Microsoft)

New hardware platforms, e.g. cloud, multi-core, etc., have led to a reconsideration of database system architecture. Our Deuteronomy project separates transactional functionality from data management functionality, enabling a flexible response to exploiting new platforms. This separation requires, however, that recovery is described logically. In this paper, we extend current recovery methods to work in this logical setting. While this is straightforward in principle, performance is an issue. We show how ARIES style recovery optimizations can work for logical recovery where page information is not captured on the log. In side-by-side performance experiments using a common log, we compare logical recovery with a state-of-the art ARIES style recovery implementation and show that logical redo performance can be competitive.

Entangled Transactions

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

As the world becomes more interdependent and computing grows more collaborative, there is a need for new abstractions and tools to help users work together. We recently introduced entangled queries - a mechanism for information exchange between database queries. In this paper, we introduce entangled transactions, units of work similar to traditional transactions that however do not run in isolation, but communicate with each other via entangled queries. Supporting entangled transactions brings about many new challenges, from an abstract model to an investigation of the unique systems issues that arise during their implementation. We first introduce a novel semantic model for entangled transactions that comes with analogues of the classical ACID properties. We then discuss execution models for entangled transactions and select a concrete design motivated by application scenarios. With a prototype system that implements this design, we show experimental results that demonstrate the viability of entangled transactions in real-world application settings.

Optimistic Concurrency Control by Melding Trees

Philip Bernstein (Microsoft Research), Colin Reid (Microsoft), Ming Wu (Microsoft Research), Xinhao Yuan (Tsinghua University)

This paper describes a new optimistic concurrency control algorithm for tree-structured data called meld. Each transaction executes on a snapshot of a multiversion database and logs a record with its intended updates. Meld processes log records in log order on a cached partial-copy of the last committed state to determine whether each transaction commits. If so, it merges the transaction's updates into that state. Meld is used in the Hyder transaction system and enables Hyder to scale out without partitioning. Since meld is on the critical path of transaction execution, it must be very fast. The paper describes the meld algorithm in detail and reports on an evaluation of an implementation. It can perform over 400K update transactions per second for transactions with two operations, and 130K for transactions with eight operations.

Research 29: Web Data

Hyper-Local, Directions-Based Ranking of Places

Petros Venetis (Stanford University), Hector Gonzalez (Google), Christian Jensen (Aarhus University), Alon Halevy (Google)

Studies find that at least 20% of web queries have local intent; and the fraction of queries with local intent that originate from mobile properties may be twice as high. The emergence of standardized support for location providers in web browsers, as well as of providers of accurate locations, enables so-called hyper-local web querying where the location of a user is accurate at a much finer granularity than with IP-based positioning. This paper addresses the problem of determining the importance of points of interest, or places, in local-search results. In doing so, the paper proposes techniques that exploit logged directions queries. A query that asks for directions from a location a to a location b is taken to suggest that a user is interested in traveling to b and thus is a vote that location b is interesting. Such user generated directions queries are particularly interesting because they are numerous and contain precise locations. Specifically, the paper proposes a framework that takes a user location and a collection of near-by places as arguments, producing a ranking of the places. The framework enables a range of aspects of directions queries to be exploited for the ranking of places, including the frequency with which places have been referred to in directions queries. Next, the paper proposes an algorithm and accompanying data structures capable of ranking places in response to hyper-local web queries. Finally, an empirical study with very large directions query logs offers insight into the potential of directions queries for the ranking of places and suggests that the proposed algorithm is suitable for use in real web search engines.

Optimal Schemes for Robust Web Extraction

Aditya Parameswaran (Stanford University), Nilesh Dalvi (Yahoo! Research), Hector Garcia-Molina (Stanford University), Rajeev Rastogi (Yahoo! India)

In this paper, we consider the problem of constructing wrappers for web information extraction that are robust to changes in websites. We consider two models to study robustness formally: the adversarial model, where we look at the worst-case robustness of wrappers, and probabilistic model, where we look at the expected robustness of wrappers, as web-pages evolve. Under both models, we present optimal algorithms for constructing the most robust wrapper. By evaluating on real websites, we demonstrate that in practice, our algorithms are highly effective in coping up with changes in websites, and reduce the wrapper breakage by up to 500% over existing techniques.

OXPath: A Language for Scalable, Memory-efficient Data Extraction from Web Applications

Tim Furche (Oxford University), Georg Gottlob (Oxford University), Giovanni Grasso (Oxford University), Christian Schallhart (Oxford University), Andrew Sellers (Oxford University)

The evolution of the web has outpaced itself: The growing wealth of information and the increasing sophistication of interfaces necessitate automated processing. Web automation and extraction technologies have been overwhelmed by this very growth. To address this trend, we identify four key requirements of web extraction: (1) Interact with sophisticated web application interfaces, (2) Precisely capture the relevant data for most web extraction tasks, (3) Scale with the number of visited pages, and (4) Readily embed into existing web technologies. We introduce OXPath, an extension of XPath for interacting with web applications and for extracting information thus revealed. It addresses all the above requirements. OXPath's page-at-a-time evaluation guarantees memory use independent of the number of visited pages, yet remains polynomial in time. We validate experimentally the theoretical complexity and demonstrate that its evaluation is dominated by the page rendering of the underlying browser. Our experiments show that OXPath outperforms existing commercial and academic data extraction tools by a wide margin. OXPath is available under an open source license.

Randomized Generalization for Aggregate Suppression over Hidden Web Databases Xin Jin (George Washington University), Nan Zhang (George Washington University), Aditya Mone (UT Arlington), Gautam Das (UT Arlington)

Many web databases are hidden behind restrictive form-like interfaces which allow users to execute search queries over the underlying hidden database. While it is important to support such search queries, many hidden database owners also want to maintain a certain level of privacy for aggregate information over their databases, for reasons including business secrets and homeland security. Existing work on aggregate suppression thwarts the uniform random sampling of a hidden database, but cannot address recently proposed attacks which accurately estimate SUM and COUNT queries without the need to first draw a uniform random sample. In this paper, we consider the problem of suppressing SUM and COUNT queries over a hidden database. In particular, we develop randomized generalization, a novel technique which provides rigid aggregate-suppression guarantee while maintaining the utility of individual search queries. We present theoretical analysis and extensive experiments to illustrate the effectiveness of our approach.

Research 30: Skyline and String Matching

ZINC: Efficient Indexing for Skyline Computation

Bin Liu (National University of Singapore), Chee-Yong Chan (National University of Singapore)

We present a new indexing method named ZINC (for Z-order Indexing with Nested Code) that supports efficient skyline computation for data with both totally as well as partially ordered attribute domains. The key innovation in ZINC is based on combining the strengths of the ZBtree, which is the state-of-the-art index method for computing skylines involving totally ordered domains, with a novel, nested coding scheme that succinctly maps partial orders into total orders. An extensive performance evaluation demonstrates that ZINC significantly outperforms the state-of-the-art TSS indexing scheme for skyline queries.

QSkycube: Efficient Skycube Computation Using Point-Based Space Partitioning Jongwuk Lee (POSTECH), Seung-won Hwang (POSTECH)

Skyline queries have gained considerable attention for multi-criteria analysis of large-scale datasets. However, the skyline queries are known to return too many results for high-dimensional data. To address this problem, a skycube is introduced to efficiently provide users with multiple skylines with different strengths. For efficient skycube construction, state-of-the-art algorithms amortized redundant computation among subspace skylines, or cuboids, either (1) in a bottomup fashion with the principle of sharing result or (2) in a top-down fashion with the principle of sharing structure. However, we observed further room for optimization in both principles. This paper thus aims to design a more efficient skycube algorithm that shares multiple cuboids using more effective structures. Specifically, we first develop each principle by leveraging multiple parents and a skytree, representing recursive point-based space partitioning. We then design an efficient algorithm exploiting these principles. Experimental results demonstrate that our proposed algorithm is significantly faster than state-of-the-art skycube algorithms in extensive datasets.

A Subsequence Matching with Gaps-Range-Tolerances Framework: A Query-By-Humming Application

Alexios Kotsifakos (University of Athens), Panagiotis Papapetrou (Aalto University), Jaakko Hollmén (Aalto University), Dimitris Gunopulos (University of Athens)

We propose a novel subsequence matching framework that allows for gaps in both the query and target sequences, variable matching tolerance levels efficiently tuned for each query and target sequence, and also constrains the maximum match length. Using this framework, a space and time efficient dynamic programming method is developed: given a short query sequence and a large database, our method identifies the subsequence of the database that best matches the query, and further bounds the number of consecutive gaps in both sequences. In addition, it allows the user to constrain the minimum number of matching elements between a query and a database sequence. We show that the proposed method is highly applicable to music retrieval. Music pieces are represented by 2-dimensional time series, where each dimension holds information about the pitch and duration of each note, respectively. At runtime, the query song is transformed to the same 2-dimensional representation. We present an extensive experimental evaluation using synthetic and hummed queries on a large music database. Our method outperforms, in terms of accuracy, several DP-based subsequence matching methods - with the same time complexity - and a probabilistic model-based method.

Approximate Substring Matching over Uncertain Strings

Tingjian Ge (University of Kentucky), Zheng Li (University of Kentucky)

Text data is prevalent in life. Some of this data is uncertain and is best modeled by probability distributions. Examples include bio-logical sequence data and automatic ECG annotations, among others. Approximate substring matching over uncertain texts is largely an unexplored problem in data management. In this paper, we study this intriguing question. We propose a semantics called (k, τ)-matching queries and argue that it is more suitable in this context than a related semantics that has been proposed previously. Since uncertainty incurs considerable overhead on indexing as well as the final verification for a match, we devise techniques for both. For indexing, we propose a multilevel filtering technique based on measuring signature distance; for verification, we design two algorithms that give upper and lower bounds and significantly reduce the costs. We validate our algorithms with a systematic evaluation on two real-world datasets and some synthetic datasets.

Industrial 1: Database Systems Testing, Debugging, and Analysis

Consistent Synchronization Schemes for Workload Replay

Konstantinos Morfonios (Oracle), Romain Colle (Stanford University), Leonidas Galanis (Oracle), Supiti Buranawatanachoke (Oracle), Benoît Dageville (Oracle), Karl Dias (Oracle), Yujun Wang (Oracle)

Oracle Database Replay has been recently introduced in Oracle 11g as a novel tool to test relational database systems. It involves recording the workload running on the database server in a production system, and subsequently replaying it on the database server in a test system. A key feature of workload replay that enables realistic reproduction of a real workload is synchronization. It is a mechanism that enforces specific ordering on the replayed requests that comprise the workload. It affects the level of request concurrency and the consistency of the replay results when compared to the captured workload. In this paper, we define the class of consistent replay synchronization schemes and study, for the first time, the spectrum they cover and the tradeoffs they present. We place the only scheme proposed so far, the one implemented in Oracle 11g Release 1, within the aforementioned spectrum and show that it is coarse-grained and more restrictive than necessary, often enforcing dependencies between calls that are independent. By enforcing needless waits, it decreases the level of possible concurrency and degrades performance. To overcome these drawbacks, we identify the best scheme within the spectrum; it is finer-grained than its counterparts and strikes the right balance across different tradeoffs: it enforces a partial ordering on the replayed calls that minimizes the number of required waits and maximizes the level of concurrency, without compromising consistency of the replay results. We have implemented the new scheme in Oracle 11g Release 2. Our experiments indicate that it produces better quality replays than the pre-existing one for major classes of workload.

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

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

We consider how to monitor and debug query processing dataflows, in distributed environments such as Pig/Hadoop. Our work is motivated by a series of informal user interviews, which revealed that monitoring and debugging needs are both pressing and diverse. In response to these interviews, we created a framework for custom dataflow instrumentation, called Inspector Gadget (IG). IG makes it easy to write a wide variety of monitoring and debugging behaviors, and attaches seamlessly to an existing, unmodified dataflow environment such as Pig. We have implemented a dozen user-requested tools in Inspector Gadget, each in just a few hundred lines of Java code. The performance overhead is modest in most cases. Our Pigbased implementation of IG, called Penny, is slated for public release in mid-2011, in conjunction with the upcoming Apache Pig v0.9 release.

HIWAS: Enabling Technology for Analysis of Clinical Data in XML Documents Joshua Hui (IBM Research - Almaden), Sarah Knoop (IBM Research - Almaden), Peter Schwarz (IBM Research - Almaden)

The information contained in large collections of clinical data can be used for many valuable purposes, such as epidemiological studies, evidence-based medicine, monitoring compliance with best clinical practices, and cost-benefit analyses. However, the emerging standards for the electronic representation of clinical data, such as the Clinical Document Architecture (CDA), are very complex and new tools are required to effectively extract and utilize the information contained in these documents. In this paper, we present HIWAS, a research prototype of a new tool that creates a structural summary of a collection of XML documents, thereby enabling users to find relevant information for a specific purpose within complex XML documents. A HIWAS user can create a target model that contains just the information they need, in a simplified representation that can be queried efficiently and is compatible with existing relational business intelligence technology. By making these complex XML documents digestible with conventional tools, HIWAS lowers a key barrier to meaningful use of aggregated clinical data.

Industrial 2: Large-Scale Analytics

Bridging Two Worlds with RICE

Philipp Große (SAP), Wolfgang Lehner (TU Dresden), Thomas Weichert (SAP), Franz Färber (SAP), Wen-Syan Li (SAP)

The growing need to use large amounts of data as the basis for sophisticated business analysis conflicts with the current capabilities of statistical software systems as well as the functions provided by most modern databases. We developed two novel approaches towards a solution for this basic conflict, based on the widely-used statistical software package R and the SAP In-Memory Computing Engine (IMCE). We thereby propose an alternative data exchange mechanism with R. Instead of using standard SQL interfaces like JDBC or ODBC we introduced SQL-SHM, a shared memory-based data exchange to incorporate R's vertical data structure. Furthermore, we extended this approach to R-Op introducing R scripts equivalent to native database operations like join or aggregation within the execution plans. With the calculation engine, IMCE provides a framework to model logical execution plans and thereby offers a convenient way to use the full functionality of R via SQL interface. Moreover, this enables us to run R scripts in parallel without the necessity of extending the R interpreter itself.

Jaql: A Scripting Language for Large Scale Semistructured Data Analysis

Kevin Beyer (IBM Research - Almaden), Vuk Ercegovac (IBM Research - Almaden), Rainer Gemulla (Max-Planck-Institut für Informatik), Andrey Balmin (IBM Research - Almaden), Mohamed Eltabakh (IBM Research - Almaden), Carl-Christian Kanne (IBM Research -Almaden), Carl-Christian Kanne (IBM Research - Almaden), Fatma Ozcan (IBM Research - Almaden), Eugene Shekita (IBM Research - Almaden)

This paper describes Jaql, a declarative scripting language for analyzing large semistructured datasets in parallel using Hadoop's MapReduce framework. Jaql is currently used in IBM's InfoSphere BigInsights and Cognos Consumer Insight products. Jaql's design features are: (1) a flexible data model, (2) reusability, (3) varying levels of abstraction, and (4) scalability. Jaql's data model is inspired by JSON and can be used to represent datasets that vary from flat, relational tables to collections of semistructured documents. A Jaql script can start without any schema and evolve over time from a partial to a rigid schema. Reusability is provided through the use of higher-order functions and by packaging related functions into modules. Most Jaql scripts work at a high level of abstraction for concise specification of logical operations (e.g., join), but Jaql's notion of physical transparency also provides a lower level of abstraction if necessary. This allows users to pin down the evaluation plan of a script for greater control or

even add new operators. The Jaql compiler automatically rewrites Jaql scripts so they can run in parallel on Hadoop. In addition to describing Jaql's design, we present the results of scale-up experiments on Hadoop running Jaql scripts for intranet data analysis and log processing.

Evaluation Strategies for Top-k Queries over Memory-Resident Inverted Indexes Marcus Fontoura (Google), Vanja Josifovski (Yahoo! Research), Jinhui Liu (Yahoo! Research), Srihari Venkatesan (Yahoo! Research), Xiangfei Zhu (Yahoo! Research), Jason Zien (Yahoo! Research)

Top-k retrieval over main-memory inverted indexes is at the core of many modern applications: from large scale web search and advertising platforms, to text extenders and content management systems. In these systems, queries are evaluated using two major families of algorithms: document-at-a-time (DAAT) and term-at-a-time (TAAT). DAAT and TAAT algorithms have been studied extensively in the research literature, but mostly in disk-based settings. In this paper, we present an analysis and comparison of several DAAT and TAAT algorithms used in Yahoo!'s production platform for online advertising. The low-latency requirements of online advertising systems mandate memory-resident indexes. We compare the performance of several query evaluation algorithms using two real-world ad selection datasets and query workloads. We show how some adaptations of the original algorithms for main memory setting have yielded significant performance improvement, reducing running time and cost of serving by 60% in some cases. In these results both the original and the adapted algorithms have been evaluated over memory-resident indexes, so the improvements are algorithmic and not due to the fact that the experiments used main memory indexes.

Industrial 3: Techniques for Large-Scale Data Management

Online Expansion of Large-scale Data Warehouses

Jeffrey Cohen (EMC), John Eshleman (EMC), Brian Hagenbuch (EMC), Joy Kent (EMC), Christopher Pedrotti (EMC), Gavin Sherry (EMC), Florian Waas (EMC)

Modern data warehouses store exceedingly large amounts of data, often considered the crown jewels of an enterprise. The amount of data maintained in such data warehouses increases significantly over time - often at a continuous pace, e.g., by gathering additional data or retaining data for longer periods to derive additional business value, but occasionally also precipitously, e.g., when consolidating disparate data warehouses and Data Marts into a single database. Having to expand a data warehouse with 100's of TB of data by a substantial portion, e.g., 100% or more is a complex and disruptive maintenance operation as it typically involves some sort of dumping and reloading of data which requires substantial downtime. In this paper we describe the methodology and mechanisms we developed in Greenplum Database to expand large-scale data warehouses in an online fashion, i.e., without noticeable downtime. At the core of our approach is a set of robust and transactionally consistent primitives that enable efficient data movement. Special emphasis was put on usability and control that lets an administrator tailor the expansion process to specific operational characteristics via priorities and schedules. We present a number of experiments to quantify the impact of an on-going expansion on query workloads.

Auto-Grouping Emails for Faster E-Discovery

Sachindra Joshi (IBM Research India), Danish Contractor (IBM Research India), Kenney Ng (IBM Software Group, USA), Prasad Deshpande (IBM Research India), Thomas Hampp (IBM Software Group, Germany)

In this paper, we examine the application of various grouping techniques to help improve the efficiency and reduce the costs involved in an electronic discovery process. Specifically, we create coherent groups of email documents which characterize either a syntactic theme, a seman-

tic theme or an email thread. All such grouped documents can be reviewed together leading to a faster and more consistent review of documents. Syntactic grouping of emails is based on near duplicate detection whereas semantic grouping is based on identifying concepts in the email content using information extraction. Email thread detection is achieved using a combination of segmentation and near duplicate detection. We present experimental results on the Enron corpus that suggest that these approaches can significantly reduce the review time and show that high precision and recall in identifying the groups can be achieved. We also describe how these techniques are integrated into the IBM eDiscovery Analyzer product offering.

Web Scale Taxonomy Cleansing

Taesung Lee (POSTECH), Zhongyuan Wang (Microsoft Research Asia), Haixun Wang (Microsoft Research Asia), Seung-won Hwang (POSTECH)

Large ontologies and taxonomies are automatically harvested from web-scale data. These taxonomies tend to be huge, noisy, and contains little context. As a result, cleansing and enriching those large-scale taxonomies becomes a great challenge. A natural way to enrich a taxonomy is to map the taxonomy to existing datasets that contain rich information. In this paper, we study the problem of matching two web scale taxonomies. Besides the scale of the problem, we address the challenge that the taxonomies may not contain enough context (such as attribute values). As existing entity resolution techniques are based directly or indirectly on attribute values as context, we must explore external evidence for entity resolution. Specifically, we explore positive and negative evidence in external data sources such as the web and in other taxonomies. To integrate positive and negative evidence, we formulate the entity resolution problem as a problem of finding optimal multi-way cuts in a graph. We analyze the complexity of the problem, and propose a Monte Carlo algorithm for finding greedy cuts. We conduct extensive experiments and compare our approach with three existing methods to demonstrate the advantage of our approach.

Industrial 4: Large-Scale Distributed Systems

Tenzing - A SQL Implementation on the MapReduce Framework

Biswapesh Chattopadhyay (Google), Liang Lin (Google), Weiran Liu (Google), Sagar Mittal (Google), Prathyusha Aragonda (Google), Vera Lychagina (Google), Younghee Kwon (Google), Michael Wong (Google)

Tenzing is a query engine built on top of MapReduce for ad hoc analysis of Google data. Tenzing supports a mostly complete SQL implementation (with several extensions) combined with several key characteristics such as heterogeneity, high performance, scalability, reliability, metadata awareness, low latency, support for columnar storage and structured data, and easy extensibility. Tenzing is currently used internally at Google by 1000+ employees and serves 10000+ queries per day over 1.5 petabytes of compressed data. In this paper, we describe the architecture and implementation of Tenzing, and present benchmarks of typical analytical queries.

An Algebraic Approach for Data-Centric Scientific Workflows

Eduardo Ogasawara (COPPE/UFRJ), Jonas Dias (COPPE/UFRJ), Daniel de Oliveira (COPPE/UFRJ), Fábio Porto (LNCC), Patrick Valduriez (INRIA), Marta Mattoso (COPPE/ UFRJ)

Scientific workflows have emerged as a basic abstraction for structuring and executing scientific experiments in computational environments. In many situations, these workflows are computationally and data intensive, thus requiring execution in large-scale parallel computers. However, parallelization of scientific workflows remains low-level, ad-hoc and labor-intensive, which makes it hard to exploit optimization opportunities. To address this problem, we propose an algebraic approach (inspired by relational algebra) and a parallel execution model that enable

automatic optimization of scientific workflows. We conducted a thorough validation of our approach using both a real oil exploitation application and synthetic data scenarios. The experiments were run in Chiron, a data-centric scientific workflow engine implemented to support our algebraic approach. Our experiments demonstrate performance improvements of up to 226% compared to an ad-hoc workflow implementation.

Citrusleaf: A Real-Time NoSQL DB which Preserves ACID V. Srinivasan (Citrusleaf), Brian Bulkowski (Citrusleaf)

In this paper, we describe the Citrusleaf real-time distributed database platform that is built using the core principles of traditional database consistency and reliability while also being fast and flexible enough for use in high-performance applications like real-time bidding. In fact, Citrusleaf is unique among NoSQL databases for its ability to provide immediate consistency and ACID while still being able to consistently exceed the high performance and scalability standards required by demanding real-time applications. This paper describes how the Citrusleaf system achieves the marriage of traditional database reliability, including immediate consistency and ACID, with flexibility and operational efficiency. Citrusleaf scales linearly at extremely high throughput while keeping response time in the sub-millisecond range as demonstrated by the test results presented here. This kind of performance has enabled Citrusleaf to become the underlying component of some of the world's largest real-time bidding networks.

Challenges and Vision 1

Data Markets in the Cloud: An Opportunity for the Database Community Magdalena Balazinska (University of Washington), Bill Howe (University of Washington), Dan Suciu (University of Washington)

Cloud-computing is transforming many aspects of data management. Most recently, the cloud is seeing the emergence of digital markets for data and associated services. We observe that our community has a lot to offer in building successful cloud-based data markets. We outline some of the key challenges that such markets face and discuss the associated research problems that our community can help solve.

Data is Dead... Without What-if Models

Peter Haas (IBM Research - Almaden), Paul Maglio (IBM Research - Almaden), Patricia Selinger (IBM Research - Almaden), Wang-Chiew Tan (IBM Research - Almaden & UCSC)

Current database technology has raised the art of scalable DESCRIPTIVE analytics to a very high level. Unfortunately, what enterprises really need is PRESCRIPTIVE analytics to identify optimal business, policy, investment, and engineering decisions in the face of uncertainty. Such analytics, in turn, rest on deep PREDICTIVE analytics that go beyond mere statistical forecasting and are imbued with an understanding of the fundamental mechanisms that govern a system's behavior, allowing what-if analyses. The database community needs to put what-if models and data on equal footing, developing systems that use both data and models to make sense of rich, real-world complexity and to support real-world decision-making. This model-and-data orientation requires significant extensions of many database technologies, such as data integration, query optimization and processing, and collaborative analytics. In this paper, we argue that data without what-if modeling may be the database community's past, but data with what-if modeling must be its future.

Data Generation for Application-Specific Benchmarking Y.C. Tay (National University of Singapore)

The Transaction Processing Council (TPC) has played a pivotal role in the growth of the database industry over the last twenty-five years. However, its handful of domain-specific benchmarks are increasingly irrelevant to the multitude of data-centric applications, and its top-down process is too slow. This mismatch calls for a paradigm shift to a bottom-up community effort to develop tools for application-specific benchmarking. Such a development program would center around techniques for synthetically scaling (up or down) an empirical dataset. This engineering effort in turn requires the development of a database theory on attribute value correlation.

Reverse Data Management

Alexandra Meliou (University of Washington), Wolfgang Gatterbauer (University of Washington), Dan Suciu (University of Washington)

Database research mainly focuses on forward-moving data flows: source data is subjected to transformations and evolves through queries, aggregations, and view definitions to form a new target instance, possibly with a different schema. This Forward Paradigm underpins most data management tasks today, such as querying, data integration, data mining, etc. We contrast this forward processing with Reverse Data Management (RDM), where the action needs to be performed on the input data, on behalf of desired outcomes in the output data. Some data management tasks already fall under this paradigm, for example updates through views, data generation, data cleaning and repair. RDM is, by necessity, conceptually more difficult to define, and computationally harder to achieve. Today, however, as increasingly more of the available data is derived from other data, there is an increased need to be able to modify the input in order to achieve a desired effect on the output, motivating a systematic study of RDM. We define the Reverse Data Management problem, and classify RDM problems into four categories. We illustrate known examples of RDM problems and classify them under these categories, and introduce a new type of RDM problem, How-To Queries.

Anthropocentric Data Systems

Peter Triantafillou (University of Patras)

We present a vision for Anthropocentric Data Systems (ADS). First, ADS is about research that develops a conduit that imports the input/wisdom of humans into the internals of data systems. Second, ADS is about research which discerns those tasks for which collective human input can actually improve system internals. Third, ADS is about researching the system architecture and structuring principles that will enable this fusion of automation and human feedback into the appropriate system internals. Finally, ADS is concerned with modeling human behavior and input, based on which such feedback can be anticipated, appropriately evaluated, and best exploited, as well as with appropriate interface systems, necessary to engage, visualize, and structure human input. We motivate and describe ADS and overview some key research challenges, including architectural principles, managing and organizing feedback (dealing with trust, maliciousness, and incentives), feedback aggregation and classification, modeling issues associating feedback to system components, prediction crowdsourcing, etc. In this, we attempt to highlight the relevant cross-disciplinary research. We also overview ongoing relevant research of our own that adds concreteness to the vision. Given the unprecedented levels of human involvement into data systems recently and the various relevant R&D results that are rapidly emerging from several different communities, the time is ripe to develop the next paradigm for creating, retrieving, and managing data, whereby collective human wisdom plays an instrumental role.

Challenges and Vision 2

Resiliency-Aware Data Management

Matthias Boehm (TU Dresden), Wolfgang Lehner (TU Dresden), Christof Fetzer (TU Dresden)

Computing architectures change towards massively parallel environments with increasing numbers of heterogeneous components. The large scale in combination with decreasing feature sizes leads to dramatically increasing error rates. The heterogeneity further leads to new error types. Techniques for ensuring resiliency in terms of robustness regarding these errors are typically applied at hardware abstraction and operating system levels. However, as errors become the normal case, we observe increasing costs in terms of computation overhead for ensuring robustness. In this paper, we argue that ensuring resiliency on the data management level can reduce the required overhead by exploiting context knowledge of query processing and data storage. Apart from reacting on already detected errors, this was mostly neglected in database research so far. We therefore give a broad overview of the background of resilient computing and existing techniques from the database perspective. Based on the lack of existing techniques on data management level, we raise three fundamental challenges of resiliency-aware data management and present example use cases. Finally, our vision of resiliency-aware data management opens many directions of future work. Fundamental research, including the partial reuse of underlying mechanisms, would allow data management systems to cope with future hardware characteristics by effectively and efficiently ensuring resiliency.

Guided Interaction: Rethinking the Query-Result Paradigm Arnab Nandi (University of Michigan), H.V. Jagadish (University of Michigan)

Many decades of research, coupled with continuous increases in computing power, have enabled highly efficient execution of queries on large databases. In consequence, for many databases, far more time is spent by users formulating queries than by the system evaluating them. It stands to reason that, looking at the overall query experience we provide users, we should pay attention to how we can assist users in the holistic process of obtaining the information they desire from the database, and not just the constituent activity of efficiently generating a result given a complete precise query. In this paper, we examine the conventional query-result paradigm employed by databases and demonstrate challenges encountered when following this paradigm for an example information seeking task. We recognize that the process of query specification itself is a major stumbling block. With current computational abilities, we are at a point where we can make use of the data in the database to aid in this process. To this end, we propose a new paradigm, guided interaction, to solve the noted challenges, by using interaction to guide the user through the query formulation, query execution and result examination processes. The user can be given advance information during query specification that cannot only assist in query formulation, but may also lead to abandonment of an unproductive query direction or the satisfaction of information need even before the query specification is complete. There are significant engineering challenges to constructing the system we envision, and the building blocks to address these challenges exist today.

The Researcher's Guide to the Data Deluge: Querying a Scientific Database in Just a Few Seconds

Martin Kersten (CWI), Stratos Idreos (CWI), Stefan Manegold (CWI), Erietta Liarou (CWI)

There is a clear need for interactive exploration of extremely large databases, especially in the area of scientific data management where ingestion of multiple Terabytes on a daily basis is foreseen. Unfortunately, current data management technology is not well-suited for such overwhelming demands. In light of these challenges, we should rethink some of the strict requirements database systems adopted in the past. We envision that next generation database

systems should interpret queries by their intent, rather than as a contract carved in stone for complete and correct answers. The result set should aid the user in understanding the database's content and provide guidance to continue the data exploration journey. A scientist can stepwise explore deeper and deeper into the database, and stop when the result quality drops below his satisfaction point. At the same time, response times should be close to instant such that they allow a scientist to interact with the system and explore the data in a contextualized way. Several research directions are carved out to realize this vision. They range from engineering a novel database kernel where speed rather than completeness is the first class citizen, up to refusing to process a costly query in the first place, but providing advice on how to reformulate it instead, or even providing alternatives the system believes might be relevant for the exploration patterns observed.

Exploring the Coming Repositories of Reproducible Experiments: Challenges and Opportunities

Juliana Freire (University of Utah), Philippe Bonnet (University of Copenhagen, Denmark), Dennis Shasha (New York University)

Computational reproducibility efforts in many communities will soon give rise to validated software and data repositories of high quality. A scientist in a field may want to query the components of such repositories to build new software workflows, perhaps after adding the scientist's own algorithms. This paper explores research challenges necessary to achieving this goal.

Databases will Visualize Queries too

Wolfgang Gatterbauer (University of Washington)

Visual Query Languages study ways to help users compose queries with visual metaphors. Information Visualization studies automatic visualization techniques to help users understand and analyze data. Query Management focuses on ways to help users manage and re-use existing queries. We observe that there is a related research question across those three topics which has not received much attention, namely that of Query Visualization: How to visually represent a query to help users quickly understand its intent? Here we argue that the involved challenges are still markedly different from those of the other three, that a solution can considerably improve the usability of DBMSs, and that the topic is thus worthy of attention. We envision, that in a few years, there will be free, modular, and lightweight tools available that allow users to visualize and interpret their queries.

Demonstration Session A: Information Integration and Information Retrieval

CerFix: A System for Cleaning Data with Certain Fixes

Wenfei Fan, University of Edinburgh & Bell Labs; Jianzhong Li, Harbin Institute of Technology; Shuai Ma, Beihang University; Nan Tangh, University of Edinburgh; Wenyuan Yu, University of Ednburgh

We present CerFix, a data cleaning system that finds cer- tain fixes for tuples at the point of data entry, i.e., fixes that are guaranteed correct. It is based on master data, editing rules and certain regions. Given some attributes of an in- put tuple that are validated (assured correct), editing rules tell us what other attributes to fix and how to correct them with master data. A certain region is a set of attributes that, if validated, warrant a certain fix for the entire tuple. We demonstrate the following facilities provided by Cer- Fix: (1) a region finder to identify certain regions; (2) a data monitor to find certain fixes for input tuples, by guiding users to validate a minimal number of attributes; and (3) an auditing module to show what attributes are fixed and where the correct values come from.

Debugging Data Exchange with Vagabond

Boris Glavic, University of Toronto; Jiang Du, University of Toronto; Renee J. Miller, University of Toronto; Gustavo Alonso, ETH Zurich; Laura M. Haas, IBM Research - Almaden

In this paper, we present Vagabond, a system that uses a novel holistic approach to help users to understand and debug data exchange scenarios. Developing such a scenario is a complex and labor-intensive process where errors are often only revealed in the target instance produced as the result of this process. This makes it very hard to debug such scenarios, especially for non-power users. Vagabond aides a user in debugging by automatically generating possible explanations for target instance errors identified by the user.

HOMES: A Higher-Order Mapping Evaluation System

Huy Vu, Oxford University; Michael Benedikt, Oxford University

We describe a system that integrates querying and query transformation in a single higherorder query language. The system allows users to write queries that integrate and combine query transformations. The power of higher-order functions also allows one to succinctly write complex relational queries. Our demonstration shows the utility of the system, explains the implementation architecture on top of a relational DBMS, and explains optimizations that combine subquery caching techniques from relational databases with sharing detection schemes from functional programming.

EIRENE: Interactive Design and Refinement of Schema Mappings via Data Examples Bogdan Alexe, UC Santa Cruz; Balder ten Cate, UCSC; Phokion Kolaitis, UCSC & IBM Research - Almaden; Wang-Chiew Tan, IBM Research - Almaden & UCSC

One of the first steps in the process of integrating information from multiple sources into a desired target format is to specify the relationships, called schema mappings, between the source schemas and the target schema. In this demonstration, we showcase a new methodology for designing these schema mappings. Our system Eirene interactively solicits data examples from the mapping designer in order to design a schema mapping between a source schema and a target schema. A data example, in this context, is a pair consisting of a source instance and a target instance showing the desired outcome of performing data exchange using the schema mapping being designed. One of the central parts of the system is a module that, given a set of such data examples, either returns a "best" fitting schema mapping, or reports that no fitting schema mapping exists.

++Spicy: an Open-Source Tool for Second-Generation Schema Mapping and Data Exchange

Bruno Marnette, INRIA Saclay & ENS Cachan; Giansalvatore Mecca, Università della Basilicata; Paolo Papotti, Università Roma Tre; Salvatore Raunich, University of Leipzig; Donatello Santoro, Università della Basilicata

Recent results in schema-mapping and data-exchange research may be considered the starting point for a new generation of systems, capable of dealing with a significantly larger class of applications. In this paper we demonstrate the first of these second-generation systems, called ++Spicy. We introduce a number of scenarios from a variety of data management tasks, such as data fusion, data cleaning, and ETL, and show how, based on the system, schema mappings and data exchange techniques can be very effectively applied to these contexts. We compare ++Spicy to the previous generations of tools, to show that this is much-needed advancement in the field.

Microsoft Codename "Montego" - Data Import, Transformation, and Publication for Information Workers

Stephen Maine, Microsoft Corporation; Lorenz Prem, Microsoft Corporation; Clemens Szyperski, Microsoft Corporation; James Terwilliger, Microsoft Corporation

One of the fundamental problems in database systems today is deriving useful information from the untold quantities of data fragments that exist in the web's data stores. Data is abundant, useful information is rare. This problem space plays host to many successful and innovative solutions from industry, and the open-source community. Each solution has its strengths and weaknesses based upon their balance between utility and usability. In this paper, we demonstrate the innovation and unique approach to data mashups that Microsoft Codename "Montego" brings to the space. The "Montego" tool allows non-technical users to create complex data queries in a graphical environment they are familiar with, while making the full expressiveness of a query language available to professional users. "Montego" operates both as a standalone client, where a user can launch it from an application like Excel* to import and manipulate data into a spreadsheet, or in as a cloud service, where a user can take the product of data transformation and publish its results into a database or to the web as an OData feed.

BROAD: Diversified Keyword Search in Databases

Feng Zhao, NUS; Xiaolong Zhang, Zhejiang University; Anthony Tung, National University of Singapore; Gang Chen, Zhejiang University

Keyword search in databases has received a lot of attention in the database community as it is an effective approach for querying a database without knowing its underlying schema. However, keyword search queries often return too many results. One standard solution is to rank results such that the "best" results appear first. Still, this approach can suffer from redundancy problem where many high ranking results are in fact coming from the same part of the database and results in other parts of the database are missed completely. In this demo, we propose the BROAD system which allows users to perform diverse, hierarchical browsing on keyword search results. Our system partitions the answer trees in the keyword search results by selecting k diverse representatives from the trees, separating the answer trees into k groups based on their similarity to the representatives and then recursively applying the partitioning for each group. By constructing summarized result for the answer trees in each of the k groups, we provide a way for users to quickly locate the results that they desire.

DivDB: A System for Diversifying Query Results

Marcos Vieira, UCR; Humberto Razente, UFABC; Maria Camila Barioni, UFABC; Marios Hadjieleftheriou, AT&T Labs - Research; Divesh Srivastava, AT&T Labs; Caetano Traina Jr., ICMC-USP; Vassilis Tsotras, UCR

With the availability of very large databases, an exploratory query can easily lead to a vast

answer set, typically based on an answer's relevance (i.e., top-k, tf-idf) to the user query. Navigating through such an answer set requires huge effort and users give up after perusing through the first few answers, thus some interesting answers hidden further down the answer set can easily be missed. An approach to address this problem is to present the user with the most diverse among the answers based on some diversity criterion. In this demonstration we present DivDB, a system we built to provide query result diversification both for advanced and novice users. For the experienced users, who may want to test the performance of existing and new algorithms, we provide an SQL-based extension to formulate queries with diversification. As for the novice users, who may be more interested in the result rather than how to tune the various algorithms' parameters, the DivDB system allows the user to provide a "hint" to the optimizer on speed vs. quality of result. Moreover, novice users can use an interface to dynamically change the tradeoff value between relevance and diversity in the result, and thus visually inspect the result as he/she interacts with this parameter. This is a great feature to the end user because finding a good tradeoff value is a very hard task and it depends on several variables (i.e., query parameters, evaluation algorithms, and dataset properties). In this demonstration we show a study of the DivDB system with two image databases that contain many images of the same object under different settings (e.g., different camera angle). We show how the DivDB helps users to iteratively inspect diversification in the query result, without the need to know how to tune the many different parameters of the several existing algorithms in the DivDB system.

FuDoCS: A Web Service Composition System Based on Fuzzy Dominance for Preference Query Answering

Karim Benouaret, LIRIS/University of Lyon; Djamal Benslimane, University of Lyon; Allel Hadjali, University of Rennes; Mahmoud Barhamgi, University of Lyon

DaaS Web service composition is a powerful means to answer users' complex queries. User preferences are a key aspect that must be taken into account in the composition scheme. In this paper, we present an approach to automatically compose DaaS Web services while taking into account the user preferences. User preferences are modeled using fuzzy sets. We use an RDF query rewriting algorithm to determine the relevant services. The fuzzy constraints of the relevant services are matched to those of the query using a set of matching methods. We rank-order services using a fuzzification of Pareto dominance, then compute the top-k service compositions. We propose also a method to improve the diversity of returned compositions while maintaining as possible the compositions with the highest scores.

AIDA: An Online Tool for Accurate Disambiguation of Named Entities in Text and Tables

Mohamed Amir Yosef, Max-Planck-Institut für Informatik; Johannes Hoffart, Max-Planck-Institut für Informatik; Ilaria Bordino, Yahoo! Research; Marc Spaniol, Max-Planck-Institut für Informatik; Gerhard Weikum, Max-Planck-Institut für Informatik

We present AIDA, a framework for online tool for entity detection and disambiguation. Given a natural-language text or a Web table, we map mentions of ambiguous names onto canonical entities like people or places, registered in a knowledge base like DBpedia, Freebase, or YAGO. AIDA is a robust framework centred around collective disambiguation exploiting the prominence of entities, similarity between the context of the mention and its candidates, and the coherence among candidate entities for all mentions. We have developed a Web-based online interface for AIDA where different formats of inputs can be processed on the fly, returning proper entities and showing intermediate steps of the disambiguation process.

Demonstration Session B: Modern Hardware, Streaming, and Benchmarking

TrustedDB: A Trusted Hardware based Outsourced Database Engine Sumeet Bajaj, Stony Brook University; Radu Sion, Stony Brook University

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

IPL-P: In-Page Logging with PCRAM

Kang-Nyeon Kim, SSungkyunkwan University; Sang-Won Lee, Sungkyunkwan University; Bongki Moon, University of Arizona; Chanik Park, Samsung Electronics; Joo-Young Hwang, Samsung Electronics Co., Ltd.

A great deal of research has been done on solid-state storage media such as flash memory and non-volatile memory in the past few years. While NAND-type flash memory is now considered a top alternative to magnetic disk drives, different types of non-volatile memory have also begun to appear in the market recently. Although some advocates of storage class memory (SCM) predicted that flash memory would give way to SCM in the very near future, we believe that they will co-exist, complementing each other, for a while until the hurdles in its manufacturing process are lifted and storage class memory becomes commercially competitive in both capacity and price. This demo presents an improved design of In-Page Logging (IPL) by augmenting it with Phase Change RAM (PCRAM) in its log area. IPL is a buffer and storage management strategy that has been proposed for flash memory database systems. Due to the byte-addressability of PCRAM and its faster speed for small reads and writes, the IPL scheme with PCRAM can improve the performance of flash memory database systems even further by storing frequent log records in PCRAM instead of flash memory. We report a few advantages of this new design that will make IPL more suitable for flash memory database systems.

HyPer-sonic Combined Transaction AND Query Processing

Florian Funke, Technische Universität München; Alfons Kemper, Technische Universität München; Thomas Neumann, Technische Universität München

In this demo we will prove that it is — against common belief — indeed possible to build a main-memory database system that achieves world-record transaction processing throughput and best-of-breed OLAP query response times in one system in parallel on the same database state. The two workloads of online transaction processing (OLTP) and online analytical processing (OLAP) present different challenges for database architectures. Currently, users with high rates of mission-critical transactions have split their data into two separate systems, one database for OLTP and one so-called data warehouse for OLAP. While allowing for decent transaction rates, this separation has many disadvantages including data freshness issues due to the delay caused by only periodically initiating the Extract Transform Load-data staging and excessive resource consumption due to maintaining two separate information systems. We present an efficient hybrid system, called HyPer, that can handle both OLTP and OLAP simultaneously by using hardware-assisted replication mechanisms to maintain consistent snapshots of the transactional data. HyPer is a main-memory database system that guarantees the full ACID properties for OLTP transactions and executes OLAP query sessions (multiple queries) on arbitrarily current and consistent snapshots. The utilization of the processor-inherent

support for virtual memory management (address translation, caching, copy-on-write) yields both at the same time: unprecedentedly high transaction rates as high as 100,000+ transactions per second and very fast OLAP query response times on a single system executing both workloads in parallel. The performance analysis is based on a combined TPC-C and TPC-H benchmark.

Proactive Detection and Repair of Data Corruption: Towards a Hassle-free Declarative Approach with Amulet

Nedyalko Borisov, Duke University; Shivnath Babu, Duke University

Occasional corruption of stored data is an unfortunate byproduct of the complexity of modern systems. Hardware errors, software bugs, and mistakes by human administrators can corrupt important sources of data. The dominant practice to deal with data corruption today involves administrators writing ad hoc scripts that run data-integrity tests at the application, database, file-system, and storage levels. This manual approach, apart from being tedious and error-prone, provides no understanding of the potential system unavailability and data loss if a corruption were to occur. We have developed the Amulet system that can verify the correctness of stored data proactively and continuously. This demonstration focuses on the uses of Amulet and its technical innovations: (i) a declarative language for administrators to specify their objectives regarding the detection and repair of data corruption; (ii) optimization and execution algorithms to meet the administrator's objectives robustly and with least cost using pay-as-you-go cloud resources; and (iii) timely notification when corruption is detected, allowing proactive repair of corruption before it impacts users and applications.

A Demonstration of HYRISE - A Main Memory Hybrid Storage Engine Martin Grund, Hasso-Plattner-Institut; Philippe Cudre-Mauroux, MIT CSAIL; Samuel Madden, MIT

We propose to demonstrate HYRISE, a main memory hybrid database system, which automatically partitions tables into vertical partitions consisting of variable numbers of columns based on access patterns to each table. Using an accurate model of cache misses, HYRISE is able to predict the performance of different partitionings, and to automatically select the best partitions using an automated database partitioning algorithm. Our demonstration will show the results of the physical partitioning based on different query workloads, allowing demo attendees to visualize, fine-tune, and modify the partitioning using a GUI. It will then show how the various physical designs affect the query plans and the performance of the database as a whole. Attendees can thus experiment with various physical models, and can grasp the potential of hybrid partitionings, which achieve a 20% to 400% performance improvement over pure all-column or all-row designs on our realistic hybrid workload derived from customer applications.

Analytics for the Real-Time Web

Maxim Grinev, ETH Zurich; Maria Grineva, ETH Zurich; Martin Hentschel, ETH Zurich; Donald Kossmann, ETH Zurich

With the emergence of mobile devices constantly connected to the Internet, the nature of usergenerated data has changed on most Web 2.0 sites. Today, people produce and share data more often and the lifespan of the data is shorter. Analyzing this data leads to new requirements for analytical systems: real-time processing and database-intensive workloads. Driven by these requirements, we have developed a new system for real-time analytics. Our system extends a key-value store, Cassandra, with push-based processing, transactional task execution, and synchronization. To demonstrate our system, we have built a service to reorganize news sites using real-time feedback from social media.

UpStream: A Storage-centric Load Management System for Real-time Update Streams Alexandru Moga, ETH Zurich; Nesime Tatbul, ETH Zurich

UpStream is a framework for load management over data streams with update semantics. It provides a novel storage manager architecture that can be plugged into data stream processing engines for serving streaming applications that require low-staleness results over real-time continuous queries. We propose to demonstrate the key aspects of the UpStream architecture as well as its performance using two different application scenarios: One that models a continuously updating financial market dashboard, and another one that is based on an intelligent transportation system for monitoring moving vehicles on a road traffic network. The demonstration will illustrate how UpStream can provide low-staleness query results for these applications under highly overloaded situations, by using a number of update scheduling and storage management techniques. This will be done through a number of interactive visual monitoring tools for the application interface as well as for monitoring the run-time operation of the UpStream system itself.

DataSynth: Generating Synthetic Data using Declarative Constraints Arvind Arasu, Microsoft Research; Raghav Kaushik, Microsoft Research; Jian Li, University of Maryland

A variety of scenarios such as database system and application testing, data masking, and benchmarking require synthetic database instances, often having complex data characteristics. We present DataSynth, a flexible tool for generating synthetic databases. DataSynth uses a simple and powerful declarative abstraction based on cardinality constraints to specify data characteristics, and uses sophisticated algorithms to efficiently generate database instances satisfying the specified characteristics. The demo will showcase various features of DataSynth using two real-world data generation scenarios.

Automatic Workload Driven Index Defragmentation

Vivek Narasayya, Microsoft Research; Hyunjung Park, Stanford University; Manoj Syamala, Microsoft Research

Queries that scan a B-Tree index can suffer significant I/O performance degradation due to index fragmentation. The task of determining if an index should be defragmented is challenging for database administrators (DBAs) since today's database engines offer no support for quantifying the impact of defragmenting an index on query I/O performance. Furthermore, DBMSs only support defragmentation at the granularity of an entire B-Tree, which can be very restrictive since defragmentation is an expensive operation and workloads typically access different ranges of an index non-uniformly. We have developed techniques to address the above two challenges, and implemented a prototype of automatic workload driven index defragmentation functionality in Microsoft SQL Server. We demonstrate this functionality by showing (a) how the system tracks the potential benefit of defragmenting an index on I/O performance at low overhead, (b) the ability to defragment a range of a B-Tree index online, and (c) how the cost/benefit trade-off can be controlled in a policy driven manner to enable automatic workload driven index defragmentation requiring minimal DBA intervention.

Demonstration Session C: MapReduce, Crowdsourcing, and Mining

Whodunit: An Auditing Tool for Detecting Data Breaches Raghav Kaushik, Microsoft Research; Ravi Ramamurthy, Microsoft Research

Commercial database systems provide support to maintain an audit trail that can be analyzed offline to identify potential threats to data security. We present a tool that performs data auditing that asks for an audit trail of all users and queries that referenced sensitive data, for example "find all queries and corresponding users that referenced John Doe's salary in the last six months". Our tool: (1) handles complex SQL queries including constructs such as grouping, aggregation and subqueries, (2) has privacy guarantees, and (3) incorporates novel optimization techniques for efficiently auditing a large workload of complex SQL queries.

CrowdDB: Query Processing with the VLDB Crowd

Amber Feng, UC Berkeley; Michael Franklin, UC Berkeley; Donald Kossmann, ETH Zurich; Tim Kraska, UC Berkeley; Samuel Madden, MIT; Sukriti Ramesh, Systems Group, ETH Zurich; Andrew Wang, UC Berkeley; Reynold Xin, UC Berkeley

Databases often give incorrect answers when data are missing or semantic understanding of the data is required. Processing such queries requires human input for providing the missing information, for performing computationally difficult functions, and for matching, ranking, or aggregating results based on fuzzy criteria. In this demo we present CrowdDB, a hybrid database system that automatically uses crowdsourcing to integrate human input for processing queries that a normal database system cannot answer. CrowdDB uses SQL both as a language to ask complex queries and as a way to model data stored electronically and provided by human input. Furthermore, queries are automatically compiled and optimized. Special operators provide user interfaces in order to integrate and cleanse human input. Currently CrowdDB supports two crowdsourcing platforms: Amazon Mechanical Turk and our own mobile phone platform. During the demo, the mobile platform will allow the VLDB crowd to participate as workers and help answer otherwise impossible queries.

InfoNetOLAPer: Integrating InfoNetWarehouse and InfoNetCube with InfoNetOLAP

Chuan Li, UIC; Philip Yu, University of Illinois at Chicago; Lei Zhao, University of Science and Technology of China; Yan Xie, University of Illinois at Chicago; Wangqun Lin, University of Illinois at Chicago

To support efficient graph OLAP operations on information networks, we propose two significant intermediate infrastructures: InfoNetWarehouse and InfoNetCube. InfoNetWarehouse is designed with novelty to be the warehouse model for information networks, which provides topic-oriented, integrated, and multi-dimensional organizational solutions for Information networks. InfoNetCube is our proposed datacube implemention that serves the OLAP of information networks. We further integrate the two infrastructures with InfoNetOLAP module into a prototype called InfoNetOLAPer, which has the following noteworthy features: (1) The basic InfoNetWarehouse schema is well implemented based on SQL Server 2000, (2) InfoNetCube improves the efficiency of InfoNetOLAP by the pre-computation of InfoNetLattice, and (3) InfoNetOLAPer supports efficient I-OLAP and T-OLAP operations.

RAMP: A System for Capturing and Tracing Provenance in MapReduce Workflows Hyunjung Park, Stanford University; Robert Ikeda, Stanford University; Jennifer Widom, Stanford University

RAMP (Reduce And Map Provenance) is an extension to Hadoop that supports provenance capture and tracing for workflows of MapReduce jobs. RAMP uses a wrapper-based approach, requiring little if any user intervention in most cases, while retaining Hadoop's parallel execution and fault tolerance. We demonstrate RAMP on a real-world MapReduce workflow generated from a Pig script that performs sentiment analysis over Twitter data. We show how RAMP's automatic provenance capture and tracing capabilities provide a convenient and efficient means of drilling-down and verifying output elements.

From SPARQL to MapReduce: The Journey Using a Nested TripleGroup Algebra HyeongSik Kim, North Carolina State University; Padmashree Ravindra, North Carolina State University; Kemafor Anyanwu, North Carolina State University

MapReduce-based data processing platforms offer a promising approach for cost-effective and Web-scale processing of Semantic Web data. However, one major challenge is that this computational paradigm leads to high I/O and communication costs when processing tasks with several join operations typical in SPARQL queries. The goal of this demonstration is to show how a system RAPID+, an extension of Apache Pig, enables more efficient SPARQL query processing on MapReduce using an alternative query algebra called the Nested TripleGroup Algebra (NTGA). The demonstration will offer opportunities for users to explore NTGA-Hadoop query plans for different SPARQL query structures as well as explore relationships between query plans based on relational algebra operators and those using NTGA operators.

MapReduce Programming and Cost-based Optimization? Crossing this Chasm with Starfish

Herodotos Herodotou, Duke University; Fei Dong, Duke University; Shivnath Babu, Duke University

MapReduce has emerged as a viable competitor to database systems in big data analytics. MapReduce programs are being written for a wide variety of application domains including business data processing, text analysis, natural language processing, Web graph and social network analysis, and computational science. However, MapReduce systems lack a feature that has been key to the historical success of database systems, namely, cost-based optimization. A major challenge here is that, to the MapReduce system, a program consists of black-box map and reduce functions written in some programming language like Java, C++, Python, or Ruby. Starfish is a self-tuning system for big data analytics that includes, to our knowledge, the first Cost-based Optimizer for simple to arbitrarily complex MapReduce programs. Starfish also includes a Profiler to collect detailed statistical information from unmodified MapReduce programs, and a What-if Engine for fine-grained cost estimation. This demonstration will present the profiling, what-if analysis, and cost-based optimization of MapReduce programs in Starfish. We will show how (non-expert) users can employ the Starfish Visualizer to (a) get a deep understanding of a MapReduce program's behavior during execution, (b) ask hypothetical questions on how the program's behavior will change when parameter settings, cluster resources, or input data properties change, and (c) ultimately optimize the program.

SocialSpamGuard: A Data Mining-Based Spam Detection System for Social Media Networks

Xin Jin, UIUC; Cindy Xide Lin, UIUC; Jiebo Luo, Kodak Research Labs; Jiawei Han, UIUC

We have entered the era of social media networks represented by Facebook, Twitter, YouTube and Flickr. Internet users now spend more time on social networks than search engines. Business entities or public figures set up social networking pages to enhance direct interactions with online users. Social media systems heavily depend on users for content contribution and sharing. Information is spread across social networks quickly and effectively. However, at the same time social media networks become susceptible to different types of unwanted and malicious spammer or hacker actions. There is a crucial need in the society and industry for security solution in social media. In this demo, we propose SocialSpamGuard, a scalable and online social media spam detection system based on data mining for social network security. We employ our GAD clustering algorithm for large scale clustering and integrate it with the designed active learning algorithm to deal with the scalability and real-time detection challenges.

GrouPeer: A System for Clustering PDMSs

Verena Kantere, Cyprus University of Technology; Dimos Bousounis, ETH Zurich; Timos Sellis, National Technical University of Athens

Sharing structured data in a PDMS is hard due to schema heterogeneity and peer autonomy. To overcome heterogeneity, peer databases employ mappings that partially match local information to that of their direct neighbors. Traditionally, a query is successively rewritten along the propagation path on each peer. This results in gradual query degradation and the inability to retrieve data pertinent to the original version, even from peers that store such data. This demonstration presents GrouPeer, a system that overcomes the query degradation problem and enables the dynamic clustering of the overlay according to the semantics of the peer data, utilizing normal query traffic. Peers are provided with a methodology that allows them to choose which rewritten version of a query to answer and discover remote information-rich sources. The demonstration illustrates the functionalities in the clustering mechanism of GrouPeer: approximate query rewriting, query similarity methodology, construction of new mappings, iterative learning process, employment of automatic schema matching, and proves the capability of the system to perform gradual semantic clustering and enable high quality answers to peer queries.

Online Visualization of Geospatial Stream Data using the WorldWide Telescope Mohamed Ali, Microsoft; Badrish Chandramouli, Microsoft Research; Jonathan Fay, Microsoft; Curtis Wong, Microsoft; Steven Drucker, Microsoft; Balan Sethu Raman, Microsoft

This demo presents the ongoing effort to meld the stream query processing capabilities of Microsoft StreamInsight with the visualization capabilities of the WorldWide Telescope. This effort provides visualization opportunities to manage, analyze, and process real-time information that is of spatio-temporal nature. The demo scenario is based on detecting, tracking and predicting interesting patterns over historical logs and real-time feeds of earthquake data.

PhD Workshop 1

Query Processing in a Self-Organized Storage System Hannes Mühleisen (Freie Universität Berlin)

The amount of scalability and robustness provided by current solutions for the storage and retrieval of data might not be sufficient to support ever-larger web applications. By using swarm intelligence to route operations in a distributed storage system, these limitations can be overcome. However, the possibilities for the efficient evaluation of complex queries in this kind of system are scarce and have not been researched yet. Based on a schema-less data model along with the building blocks for complex queries on this model, I present an approach for complex query processing as part of my PhD work. Here, complex queries are moved through the distributed storage system, while constantly being re-optimized using strictly local information. The approach is described along with an evaluation methodology and a test protocol. My goal is to contribute complex query processing for a fully distributed storage system with unreliable basic read operations and probabilistic directional content routing.

Efficient Top-k Searching According to User Preferences Based on Fuzzy Functions With Usage of Tree-Oriented Data Structures

Matus Ondreicka (Charles University in Prague)

In the last few years, research of top-k processing is in progress in various domains. The aim of our research is efficient top-k searching of the best k objects with more attributes according to more complex user preferences. We focus on a multi-user solution, where data is common for all users. We use a model of user preferences based on fuzzy functions, where each user can express his/her preferences for each attribute by a fuzzy function and mutual relations between the attributes by an aggregation function. In our previous research we focused on top-k searching in tree-oriented data structures and we developed various top-k algorithms, which solve top-k problem efficiently and support the model of user preferences. Furthermore, we assume that attributes of an object type can be distributed in various information resources, e.g. data streams, web services, remote RDBMS, and so on. Therefore, integration of data is one of the topics of our research. Our research also focuses on development of the TREETOPK software system that will be including our new-developed technologies and experiments with real data from various information resources.

Top-k Web Service Compositions in the Context of User Preferences *Karim Benouaret (University of Lyon)*

Data-driven Web service composition is a powerful means to answer users' complex queries. User preferences are a key aspect that must be taken into account in the composition scheme. In this paper, we present an approach to automatically compose Data-drive Web services while taking into account the user preferences. User preferences are modeled using fuzzy sets, then incorporated into the composition queries. We use an RDF query rewriting algorithm to determine the relevant services. The fuzzy constraints of the relevant services are matched to those of the query using a set of matching methods. We define a criterion to rank-order available services using a fuzzification of Pareto dominance, then compute the top-k service compositions. We propose also a method to improve the diversity of returned compositions while maintaining as possible the compositions with the highest ranking ones. Finally, we present the architecture of our system.

PhD Workshop 2

Mixed Workload Management for In-Memory Databases Johannes Wust (Hasso Plattner Institute)

The fast-paced business environment in most industries requires analytical applications to report on the latest data of an enterprise. Today, organizations either rely on data warehouses, which make flexible reporting on up-to-date data almost impossible, or maintain copies of operational data, leading to increased costs and architectural complexity. In-memory database management systems (DBMS) on multicore CPUs are a promising platform to run analytical queries (OLAP) directly on operational data, in addition to processing load from transactional queries (OLTP). To meet response time requirements of both types of queries on a single system, some form of workload management has to be implemented to avoid system overload. Consequently, the early-stage PhD project described in this paper targets to design a workload manager optimized for mixed workloads on in-memory databases. At this stage, we envision a DBMS-internal workload management system that assigns as many resources to OLTP queries as necessary to meet response time requirements and schedules OLAP queries on the remaining resources. To minimize overall OLAP response time, we intend to leverage the underlying multicore architecture by scheduling available resources on the level of atomic database tasks, thereby considering cache behavior and the degree of parallelism. We believe that existing solutions in this area fall short for two reasons: (1) they do not consider the different response time requirements of both, operational and analytical applications, and (2) they do not take characteristics of modern in-memory DBMS, such as cache access patterns and massive parallel execution, into account. Our findings will be verified in the context of enterprise computing by using workloads from real-life enterprise systems.

Scaling Web Applications: A Temporal Approach Zhiwu Xie (University of New Mexico)

In this paper we present the theoretical aspects of a consistent and scalable web application framework. In this framework the server handles all updates and maintains a centralized temporal database, but only answers queries used for replicating snapshot databases for pre-defined periods. These database snapshots are subsequently replicated to the client-side database to answer the general web queries. The scalability is improved because most queries don't require synchronization, and in case synchronization is required, it makes effective use of the existing web cache to reduce latency and queries on the centralized database. We present the theoretic underpinning of this approach as well as the tradeoffs imposed.

Research on a Schema and Data Versioning System Bob Wall (Montana State University)

In this paper, I describe my dissertation research. This primarily involves creating a framework for effectively managing the inevitable evolution of schema in an Online Transaction Processing (OLTP) database. This framework is based on common practices used to manage source code changes in software development. It allows users of a database to create sandboxes in which changes to the database are isolated from the main database and from other sandboxes. Schema versioning techniques are used to isolate schema changes made within sandboxes and to allow queries executed in a sandbox to retrieve data from the main database without converting the database to the schema associated with the sandbox. The framework also includes a mechanism for handling data versioning within each sandbox. The combination of these two versioning techniques provides a system that significantly reduces the effort, storage capacity, and infrastructure required to perform development, maintenance, and testing of schema and data changes against a high-volume online database.

PhD Workshop 3

Matching Tree Patterns on Partial-trees Shachar Harussi (Tel Aviv University), Amir Averbuch (Tel Aviv University)

Tree patterns are query language that are used to simplify processing while making it more efficient. The existing methodologies for tree pattern matching in database (DB) process strings of labels that annotate nodes in the tree. The existing structural-join and structural-index techniques process strings of labels. My thesis focuses on developing methodologies that enable structural-join and structural-index processing on labels of nodes in subtrees. This thesis introduces a simple tree model that is called partial-tree that models both tree patterns and tree data as partialtrees. Partial-trees supply an infrastructure for efficient processing of both structural-join and structural-index on subtrees. This thesis develops two methods for pattern matching: multi-scale index and sub-linear join. Tree pattern matching in partial-trees also enables to perform query on data that is partially known such as probabilistic or experimental data. The thesis suggests an automata based approach that models partial-trees as tree automata, in a tree. It also defines how to match partial-tree patterns, which are formalized as tree automata, in partial-tree data that is also formalized as tree automata.

Knowledge-Based Complex Event Processing Kia Teymourian (Freie Universität Berlin)

Many of real-world events cannot be processed by the existing event processing systems because they are too complex to be understood and processed by the systems. Complex events can be inferred from raw primitive events based on their incoming sequence, their syntax and semantics. Usage of ontological knowledge about events and their relationship to other non-event concepts in the application domain, can improve the quality of event processing. In this Ph.D. thesis, I am aiming to address the challenges of adding formalized vocabularies/ontologies and declarative rules to the area of event processing for enabling more intelligent event processors which can understand the meaning of events and are enable to trigger meaningful reactions.

3rd TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2011)

Workshop Organizers: Raghunath Nambiar, Meikel Poess Location: Cascade 1BC

Monday 8:30 AM - 9:15 AM

Opening Remarks and Welcome

TPC Technology Conference on Performance Evaluation and Benchmarking Raghunath Nambiar, General Chair (Cisco)

TPC Public Relations Presentation

Shaping the Landscape of Industry Standard Benchmarks: Contributing of the Transaction Processing Performance Council (TPC)

Raghunath Nambiar (Cisco); Nicholas Wakou (Dell); Peter Thawley (Sybase) , Andrew Maslan (NEC); Matthew Lanken (Oracle); Michael Majdalany (LoBue & Majdalany Magt Group); Forrest Carman (Owen Media)

Monday 9:15 AM - 10:05AM

Session 1: New Benchmark Ideas

Metrics for Measuring the Performance of the Mixed Workload CH-benCHmark Florian Funke, Alfons Kemper, Stefan Krompass, Thomas Neumann, Michael Seibold (Technische Universität München); Harumi Kuno (HP Labs); Anisoara Nica (Sybase); Meikel Poess (Oracle)

Towards an Enhanced Benchmark Advocating Energy-Effcient Systems

Daniel Schall, Volker Hoefner (University of Kaiserslautern); Manuel Kern (SPH AG)

Monday 10:30 AM - 11:10 AM

Keynote

Optimization of Analytic Data Flows for Next Generation Business Intelligence Applications

Umeshwar Dayal (HP Labs)

Monday 11:10 AM – 12:00 PM

Session 1: New Benchmark Ideas (continued)

Normalization in a Mixed OLTP and OLAP Workload Scenario

Anja Bog, Alexander Zeier (Hasso Plattner Institute); Kai Sachs (SAP AG)

Performance Metrics for Complex Event Processing Systems Torsten Grabs, Ming Lu (Microsoft)

Monday 1:30 PM – 2:10 PM

Invited Talk

Benchmarking With Your Head in the Cloud Karl Huppler (IBM)

Monday 2:10 PM – 3:55 PM

Session 2: Extending TPC Benchmarks

Invited Talk: Extending TPC-E to Measure Availability in Database Systems Yantao Li, Charles Levine (Microsoft)

A PDGF Implementation for TPC-H

Meikel Poess (Oracle); Tilmann Rabl (Universität Passau), Manuel Danisch (Universität Passau), Michael Frank (Universität Passau)

Coffee Break

Introducing Skew into the TPC-H Benchmark

Alain Crolotte, Ahmad Ghazal (Teradata)

Monday 3:55 PM – 5:35 PM

Session 3: Performance Modeling and Characterization

Time and Cost-Efficient Modeling and Generation of Large-Scale TPC Workloads Christina Delimitrou, Christos Kozyrakis (Stanford University); Sriram Sankar, Badriddine Khessib, Kushagra Vaid (Microsoft)

When free is not really free: What does it cost to run a database workload in the cloud?

Avrilia Floratou, Willis Lang, Jignesh Patel (University of Wisconsin); Alan Halverson (Microsoft)

A Fine-grained Performance-based Decision Model for Virtualization Application Solution

Jianhai Chen, Dawei Huang, Bei Wang, Deshi Ye, Qinming He, Wenzhi Chen (Zhejiang University)

SI-CV: Snapshot Isolation With Co-Located Versions Robert Gottstein, Ilia Petrov, Alejandro Buchmann (TU-Darmstadt)

Monday 5:35 PM - 6:00 PM

Closing Remarks: Benchmarking for 2012 and Beyond

Meikel Poess, General Chair (Oracle)

9th International Workshop on Quality in Databases (QDB 2011)

Workshop Organizers: Mourad Ouzzani, Paolo Papotti, Erhard Rahm Location: Vashon

Monday 9:00 AM – 10:00 AM

Keynote 1: Insight into Data Cleaning and Linkage

Ahmed K. Elmagarmid (Qatar Computing Research Institute)

Monday 10:30 AM – 12:00 PM

Research Session 1

Macro-level Scheduling of ETL Workflows

Anastasios Karagiannis, Panos Vassiliadis (University of Ioannina), Alkis Simitsis (HP Labs)

Infrastructures and Bounds for Distributed Entity Resolution

Csaba Sidlo, Andras Garzo, Andras Molnar, Andras Benczur (Hungarian Academy of Sciences)

The Case for CostSensitive and Easy-To-Interpret Models in Industrial Record Linkage

Vitor Carvalho, Andrew Borthwick (Intelius Data Research), Sheng Chen (Stevens Institute of Technology)

Monday 1:30 PM – 3:00 PM

Keynote 2: Data Quality: Theory and Practice

Wenfei Fan (University of Edinburgh)

Session 2

Using Argumentation to Support the User Involvement In Data Cleaning Emanuel Santos, Helena Galhardas (INESC-ID and Technical University of Lisbon)

Monday 3:30 PM - 4:30 PM

Session 3

Transformation Lifecycle Management with Nautilus Melanie Herschel, Torsten Grust (Universitaet Tubingen)

Discovering Pattern Tableaux for Data Quality Analysis Lukasz Golab, Flip Korn, Divesh Srivastava (AT&T Labs - Research)

8th International Workshop on Data Management for Sensor Networks (DMSN 2011)

Workshop Organizers: Mohamed Mokbel, Suman Nath, Wang-Chien Lee, Demetris Zeinalipour Location: St. Helens

Monday 8:30 AM – 10:00 AM

Keynote: Understanding the Genome of Data Centers Jie Liu (Microsoft Research)

Monday 10:30 AM – 12:00 PM

Research Session 1

Combining Sensor Selection with Routing and Scheduling in Wireless Sensor Networks

Conor Muldoon, Niki Trigoni, Gregory O Hare (University of Oxford)

On Quantifying the Accuracy of Maximum Likelihood Participant Reliability Estimation in Social Sensing

Dong Wang and Tarek Abdelzaher (University of Illinois, Urbana-Champaign)

Scalable Data Partitioning Techniques for Parallel Sliding Window Processing over Data Streams

Cagri Balkesen, Nesime Tatbul (ETH Zurich)

Monday 1:30 PM - 3:00 PM

Research Session 2

Sensor Data Stream Exploration for Monitoring Applications Shengliang Xu, Magdalena Balazinska (University of Washington)

Towards a Network-aware Wireless Sensor Middleware Architecture

Panayiotis Andreou, Demetris Zeinalipour-Yazti, George Samaras (University of Cyprus), Panos Chrysanthis (University of Pittsburgh)

On the Use of Histograms for Processing Exact Aggregate Queries in Wireless Sensor Networks

Khaled Ammar and Mario A. Nascimento (University of Alberta)

Monday 3:30 PM – 5:00 PM

Panel

DBRank 2011 and MUD 2011

5th International Workshop on Ranking in Databases (DBRank 2011)

Workshop Organizers: Kaushik Chakrabarti (Microsoft Research, USA) Davide Martinenghi (Politecnico di Milano, Italy)

5th International Workshop on Management of Uncertain Data (MUD 2011)

Workshop Organizers: Maurice van Keulen, Ander de Keijzer, Ghita Berrada Location: Cascade 2

Monday 8:30 AM – 10:00 AM

Keynote: Preferences, Queries, Logics

Jan Chomicki (University at Buffalo)

Monday 10:30 AM – 12:00 PM

Session 1

Toward Efficient SPARQL Top-K Query Answering in Virtual RDF Repositories Sara Magliacane (Politecnico di Milano), Alessandro Bozzon (Politecnico di Milano), Emanuele Della Valle (Politecnico di Milano)

Quality Impact of Value Matching and Scoring in Top-k Entity Attribute Extraction Matthew Solomon (Columbia University), Luis Gravano (Columbia University), Cong Yu (Google Research)

Efficient Evaluation of Semi-Skylines

Markus Endres (University of Augsburg), Werner Kießling (University of Augsburg)

Monday 1:30 PM – 3:00 PM

Session 2

Optimizing Candidate Document Selection for Query Answering in XML Databases

Mariam Salloum (University of California), Vassilis J. Tsotras (University of California), Divesh Srivastava (AT&T Labs-Research), Luna Dong (AT&T Labs-Research)

Parallel Data Access for Top-K Pipe Join

Adnan Abid (Politecnico di Milano), Marco Tagliasacchi (Politecnico di Milano)

Towards Large-Scale Multimedia Exploration

Christian Beecks (RWTH Aachen University), Tomas Skopal (Charles University in Prague), Klaus Schoeffmann (Klagenfurt University), Thomas Seidl (RWTH Aachen University)

Monday 3:30 PM – 5:30 PM

Session 3

BeyOND — Unleashing BOND

Thomas Bernecker (Ludwig-Maximilians-Universität München), Franz Graf (Ludwig-Maximilians-Universität München), Hans-Peter Kriegel (Ludwig-Maximilians-Universität München), Christian Moennig Ludwig-Maximilians-Universität München), Arthur Zimek (Ludwig-Maximilians-Universität München)

An Index Structure for Spatial Range Querying on Gaussian Distributions Kazuki Kodama (Nagoya University, Japan), Tingting Dong (Nagoya University, Japan), Yoshiharu Ishikawa (Nagoya University, Japan)

Named Entity Extraction and Disambiguation: The Reinforcement Effect

Mena Badieh Habib Morgan (University of Twente, The Netherlands), Maurice van Keulen (University of Twente, The Netherlands)

Fast Frequent Itemset Mining from Uncertain Databases Using GPGPU

Yusuke Kozawa (University of Tsukuba, Japan), Toshiyuki Amagasa (University of Tsukuba, Japan), Hiroyuki Kitagawa (University of Tsukuba, Japan)

13th International Symposium on Database Programming Languages (DBPL 2011)

Workshop Organizers: Nate Foster, Anastasios Kementsietsidis Location: Cascade 1A

Monday 8:30 AM - 9:30 AM

Invited Talk

Databases and Programming Languages: Together again for the first time Philip Wadler

Monday 9:30 AM - 10:00 AM

Session 1

Remote Batch Invocation for Database Access William Cook, Ben Wiedermann

Monday 10:30 AM – 12:00 PM

Session 2

- DBWiki: A Database Wiki prototyped in Links James Cheney, Sam Lindley, Heiko Müller
- Chasing One's Tail: XPath Containment Under Cyclic DTDs Peter Wood, Mahtab Montazerian

On guarded simulations and acyclic first-order languages George Fletcher, Jan Hidders, Stijn Vansummeren, Yongming Luo, Francois Picalausa, Paul De Bra

Monday 1:30 PM – 2:30 PM

Invited Talk

Programming and Debugging Large-Scale Data Processing Workflows Christopher Olston

Monday 2:30 PM - 3:00 PM

Session 3

Temporal Data Model for Program Debugging Demian Lessa, Bharat Jayaraman, Jan Chomicki

Monday 3:30 PM - 5:00 PM

Session 4

PSPARQL Query Containment

Melisachew Wudage Chekol, Jérôme Euzenat, Pierre Genevès, Nabil Layaïda

Next Generation Database Programming and Execution Environment

Dirk Habich, Matthias Boehm, Maik Thiele, Benjamin Schlegel, Ulrike Fischer, Hannes Voigt, Wolfgang Lehner

Validity of Positive XPath Queries with Wildcard in the Presence of DTDs Kenji Hashimoto, Yohei Kusunoki, Yasunori Ishihara, Toru Fujiwara

2nd International Workshop on Accelerating Data Management Systems Using Modern Processors and Storage Architectures (ADMS 2011)

Workshop Organizers: Rajesh Bordawekar, Christian Lang Location: St. Helens

Friday 9:00 AM – 10:30 AM

Welcome and Keynote Welcome

Keynote: Looking for Acceleration -- a practitioner's perspectives Sumanta Chatterjee, VP of Development (Oracle Database Group)

Friday 10:45 AM – 12:15 AM

Session 1: Optimizing for Multi-core Processors

Chair: Rajesh Bordawekar

Fast Sorted-Set Intersection using SIMD Instructions

Benjamin Schlegelm (Technische Universitat Dresden), Thomas Willhalm (Intel GmbH), Wolfgang Lehner (Technische Universitat Dresden)

Efficient XML Path Filtering Using GPUs

Roger Moussalli, Robert Halstead, Mariam Salloum, Walid Najjar, Vassilis J. Tsotras (University of California, Riverside)

Applicability of GPU Computing for Efficient Merge in In-Memory Databases

Jens Krueger, Martin Grund, Ingo Jaeckel, Alexander Zeier, Hasso Plattner (Hasso Plattner Institute, University of Potsdam)

Friday 1:45 PM - 3:15 PM

Session 2: Optimizing for Memory Sub-systems

Chair:: Rajesh Bordawekar

Accelerating Foreign-Key Joins using Asymmetric Memory Channels

Holger Pirk, Stefan Manegold, Martin Kersten (CWI)

Fast, Energy Efficient Scan inside Flash Memory

Sungchan Kim (Chonbuk National University), Hyunok Oh (Hanyang University), Chanik Park (Samsung Electronics), Sangyeun Cho (University of Pittsburgh), Sang-Won Lee (Sungkyunkwan University)

Do Query Optimizers Need to be SSD-aware?

Steven Pelley, Kristen Lefevre and Thomas F. Wenisch (University of Michigan)

Friday 3:45 PM – 5:30 PM

Panel Session: In the multi-core age, how do larger, faster, cheaper and more responsive memory sub-systems affect data management?

Confirmed panelists: Philip Bernstein (Microsoft Research), D. K. Panda (Ohio State University), Ken Ross (Columbia University), Guy Lohman (IBM Research - Almaden), Alexander Zeier (Hasso Plattner Institute)

5th Workshop on Large-Scale Distributed Systems and Middleware (LADIS 2011)

Workshop Organizers: Ymir Vigfusson, Marcos K. Aguilera, Gregory Chockler, Marc Shapiro, Antonio Sousa Location: Cascade 2

Friday 8:45 AM - 9:45 AM

Session 1: Scalable Networking (Keynote)

Chair: Divy Agrawal

Keynote: Bringing scale out growth of capacity to data center networks Amin Vahdat (Google)

Friday 10:15 AM – 12:15 PM

Session 2: Big Data Analysis

Chair: Gregory Chockler

Can MapReduce learn form Materialized Views?

Thomas Jörg, Roya Parvizi, Hu Yong and Stefan Dessloch (Technische Universität Kaiserslautern)

Microsharding: A Declarative Approach to Support Elastic OLTP Workloads Junichi Tatemura and Hakan Hacigumus (NEC Laboratories America)

GLADE: A Scalable Framework for Efficient Analytics

Florin Rusu and Alin Dobra (University of California, Merced / University of Florida)

Two Parallel Approaches to Network Data Analysis

Arian Baer, Antonio Barbuzzi, Pietro Michiardi and Fabio Ricciato (FTW / EURECOM)

Friday 1:30 PM – 3:00 PM

Session 3: Scalable Infrastructures

Chair: Marc Shapiro

SLO-Driven Right-Sizing and Resource Provisioning of MapReduce Jobs Abhishek Verma, Ludmila Cherkasova and Roy Campbell (University of Illinois / Hewlett-Packard Labs)

Distributed Indexing and Locking : In Search for Scalable Consistency

Mahdi Tayarani Najaran, Primal Wijesekera, Andrew Warfield and Norm C. Hutchinson (University of British Columbia)

Pyramid: A large-scale array-oriented active storage system Viet-Trung Tran, Bogdan Nicolae, Gabriel Antoniu and Luc Bougé (ENS Cachan / INRIA)

Friday 3:30 PM - 4:30 PM

Session 4: Scalable Storage (Keynote)

Chair: Divy Agrawal Keynote: Falcon: Transactional Storage at the Speed of Flash Dahlia Malkhi (MSR Silicon Valley)

Friday 6:00 PM – 9:30 PM

Friday night LADIS cruise Gather at 6:00 PM; Cruise: 6:30-9:30 PM (optional ticket event)

Saturday 9:00 AM – 10:00 AM

Session 5

Chair: Divy Agrawal

Keynote

Big Data Platforms Shaping the Next Generation Big Analytics in Enterprises Hamid Pirahesh (IBM Research)

Saturday 10:30 AM - 12:00 PM

Session 6: Resource Monitoring in the Cloud

Chair: Ymir Vigfusson

Diagnosing Latency in Multi-Tier Black-Box Services Krzysztof Ostrowski, Gideon Mann and Mark Sandler (Google Inc.)

Characterizing Task Usage Shapes in Google's Compute Clusters Qi Zhang, Raouf Boutaba and Joseph Hellerstein (University of Waterloo / Google Inc.)

CIRCUMFLEX: A Scheduling Optimizer for MapReduce Workloads With Shared

Scans

Joel Wolf, Andrey Balmin, Deepak Rajan, Kirsten Hildrum, Rohit Khandekar, Sujay Parekh, Kun-Lung Wu and Rares Vernica (IBM Research)

Saturday 1:30 PM – 3:00 PM

Session 7: Security and Reliability in the Cloud

Chair: Luis Rodrigues

Commensal Cuckoo: Secure Group Partitioning for Large-Scale Services Siddhartha Sen and Michael Freedman (Princeton University)

Adaptive and Dynamic Funnel Replication in Clouds

Guy Laden, Roie Melamed and Ymir Vigfusson (IBM Research / Reykjavik University)

FRAPPE: Fast Replication Platform for Elastic Services

Vita Bortnikov, Gregory Chockler, Dmitri Perelman, Alexey Roytman, Shlomit Shachor and Ilya Shnayderman (IBM Research / Technion, Israel Institute of Technology)

Saturday 3:30 PM – 4:30 PM

Session 8:

Chair: Divy Agrawal Keynote Challenging Consistency Eric Brewer (UC Berkeley, Google)

5th International Workshop on Real-Time Business Intelligence (BIRTE 2011)

Workshop Organizers: Umesh Dayal, Malu Castellanos, Wolfgang Lehner Location: Fifth Avenue

8:30 AM - 10:00 AM

Session 1: Opening & Welcome

Chair: Malu Castellanos (HP Labs) Keynote: Blink: Not Your Father's Database!

Guy Lohman (IBM Research - Almaden)

10:30 AM - 12:15 PM

Session 2: Innovative System Architectures

Chair: Sang Kyun Cha (SAP, Korea)

MemcacheSQL - A Scale-Out SQL Cache Engine

Qiming Chen, Meichun Hsu, Ren Wu (HP Labs, USA)

A Cost-Aware Strategy for Merging Differential Stores in Column-Oriented In-Memory DBMS

Florian Huebner (Hasso Plattner Institute), Joos-Hendrik Boese (SAP), Jens Krueger (Hasso Plattner Institute), Frank Renkes (SAP), Cafer Tosum (SAP), Alexander Zeier (Hasso Plattner Institute)

Microsoft SQL Server Parallel Data Warehouse – Architecture Overview (invited talk) José Blakeley (Microsoft Corporation)

1:30 PM – 3:15 PM

Session 3: Novel Query Support

Chair: Umesh Dayal (HP Labs)

SAP HANA: Breaking Vertical and Horizontal Tiers in Enterprise with High-Performance In-Memory Database (invited talk)

Sang Kyun Cha (SAP, Korea)

Relax and Let the Database do the Partitioning Online. Alekh Jindal, Jens Dittrich (Saarland University, Germany)

Adaptive Processing of Multi-Criteria Decision Support Queries

Venkatesh Raghavan , Shweta Srivastava, Venkatesh Raghavan, Elke Rundensteiner (Worcester Polytechnic Institute)

3:45 PM – 5:00 PM

Session 4: Innovative Applications

Chair: Wolfgang Lehner (TU Dresden)

Scalable Social-Graphing Analytics with the Vertica Analytic Platform (invited talk) Shilpa Lawande, Andrew Lamb, Lakshmikant Shrinivas (Hewlett-Packard)

A Near Real-Time Personalization for eCommerce Platform Amit Rustagi (eBay)

Closing Remarks

Malu Castellanos, Wolfgang Lehner

8th Workshop on Secure Data Management (SDM 2011)

Workshop Organizers: Willem Jonker, Milan Petkovic, Luan Ibraimi Location: Cascade 1A

Friday 9:00 AM - 10:00 AM

Welcome / Purpose Control

Welcome

Purpose Control: did you process the data for the intended purpose? Milan Petkovic, Davide Prandi, Nicola Zannone

Friday 10:30 AM – 12:00 PM

Privacy Protection and Quantification

Quantifying Privacy Violations

Mishtu Banerjee, Rosa Karimi Adl, Leanne Wu, Ken Barker

On-the-Fly Generalization Hierarchies for Numerical Attributes Revisited Alina Campan, Nicholas Cooper and Traian MariusTruta

LORA: Link Obfuscation by Randomization in Graphs QianXiao, Zhengkui Wang, Kian-Lee Tan

Friday 1:30 PM – 3:00 PM

Security in Cloud and Sensor Networks

A Comprehensive Framework for Secure Query Processingon Relational Data in the Cloud

Shiyuan Wang, Divyakant Agrawal, AmrEl Abbadi

Challenges in Secure Sensor-Cloud Computing Nayot Poolsappasit, Vimal Kumar, Sanjay Madria, Sriram Chellappan

SensorSafe: a Framework for Privacy-PreservingManagement of Personal Sensory Information

Haksoo Choi, Supriyo Chakraborty, Zainul M. Charbiwala, Mani B.Srivastava

Friday 3:30 PM - 5:00 PM

Secure Data Management Technologies

Watermarking for Adaptive Streaming Protocols Dmitri Jarnikov, Jeroen M. Doumen

Conjunctive Wildcard Search over Encrypted Data Christoph Bosch, Richard Brinkman, Pieter Hartel, Willem Jonker

Adjusting the Trade-off between Privacy Guaranteesand Computational Cost in Secure Hardware PIR

115

Spiridon Bakiras, Konstantinos F. Nikolopoulos

1st Workshop on Searching and Integrating New Web Data Sources (VLDS 2011)

Workshop Organizers: Marco Brambilla, Fabio Casati, Stefano Ceri Location: Vashon

Friday 8:30 AM – 9:30 AM

Session 1

Opening and Academic Keynote

Welcome and Opening

Information Seeking: Convergence of Search, Recommendations and Advertising Hector Garcia-Molina (Stanford University)

Friday 9:30 AM – 10:10 AM

Session 2

Personal and Social Search

On Principles of Egocentric Person Search in Social Networks

Sara Cohen(Hebrew University of Jerusalem), Benny Kimelfeld, Georgia Koutrika and Jan Vondrak (IBM Research – Almaden)

The Power of Integrated Abstraction for Data-Centric Human/Machine

Computations

Atsuyuki Morishima, Norihide Shinagawa and Shoji Mochizuki (University of Tsukuba)

Friday 10:40 AM – 12:00 PM

Session 3

Data Extraction and Linking

Thematic Exploration of Linked Data

Silvana Castano, Alfio Ferrara and Stefano Montanelli (Università degli Studi di Milano)

Generating Linked Data by Inferring the Semantics of Tables Varish Mulwad, Tim Finin and Anupam Joshi (University of Maryland)

Contextual Data Extraction and Instance-Based Integration

Lorenzo Blanco, Valter Crescenzi, Paolo Merialdo and Paolo Papotti (Università di Roma 3)

Growing Triples on Trees: an XML-RDF Hybrid Model for Annotated Documents François Goasdoué, Konstantinos Karanasos, Yannis Katsis, Julien Leblay, Ioana Manolescu and Stamatis Zampetakis (INRIA Saclay and LRI, Université Paris-Sud)

Friday 1:30 PM – 2:10 PM

Session 4

Industrial Keynote

Towards an Eco-System of Structured Data on the Web Alon Halevy (Google)

Friday 2:10 PM – 3:10 PM

Session 5

Using Ranking for Selection, Visualization, and Valorization

Searching Social Updates for Topic-centric Entities

Maria Christoforaki, Ivie Erunse and Cong Yu (New York University)

Enhanced Visualization for Web-Based Summaries Brent Wenerstrom and Mehmed Kantardzic (Duthie Center for Engineering, Lousville)

Designing a Revenue Mechanism for Federated Search Engines

Laura Ernesta Bonetti, Sofia Ceppi and Nicola Gatti (Politecnico di Milano)

Friday 3:30 PM - 4:30 PM

Session 6

New and News Applications

Dynamic Visualizations for Multi-Domain Search Results

Alessandro Bozzon, Luca Cioria, Piero Fraternali and Maristella Matera (Politecnico di Milano)

Search Computing Meets Data Extraction

Alessandro Bozzon, Chiara Pasini, Luca Tettamanti, Salvatore Vadacca (Politecnico di Milano), Tim Furche and Giorgio Orsi (Oxford University)

Use of Query Similarity for Improving Presentation of News Verticals

Annie Louis (University of Pennsylvania), Eric Crestan (Microsoft), Youssef Billawala, Rao Shen, Fernando Diaz and Jean-Francois Crespo (Yahoo! Labs)

Friday 4:30 PM – 5:00 PM

Session 7

Discussion on Very Large Data Search animated by the Workshop Organizers

Closing Discussion

5th Workshop on Personalized Access, Profile Management, and Context Awareness in Databases (PersDB 2011)

Workshop Organizers: Zoltan Gyongyi, Yannis Stavrakas Location: Cascade 1BC

Friday 8:30 AM – 10:00 AM

Session 1: Opening and Keynote

It's Y!ou---Content Optimization at Yahoo! Raghu Ramakrishnan (Yahoo! Research, USA)

Ragnu Ramakrishnan (Yanoo! Research, Os

Friday 10:30 AM – 12:00 PM

Session 2: Personalized, Social, and Context-aware Computing and Applications

Personalization, Socialization, and Recommendations in Location-based Services 2.0 Mohamed Mokbel (University of Minnesota), Jie Bao (University of Minnesota), Ahmed Eldawy (University of Minnesota), Justin Levandoski (University of Minnesota), Mohamed Sarwat (University of Minnesota)

Accelerating Problem Solving in Collaborative Social Networks (Invited Paper)

Yi Chen (Arizona State University, USA), Shu Tao (IBM Research, USA), Xifeng Yan (University of California at Santa Barbara, USA)

KANIS: Preserving k-Anonymity Over Distributed Data

Katerina Doka (National Technical University of Athens, Greece), Dimitrios Tsoumakos (National Technical University of Athens, Greece), Nectarios Koziris (National Technical University of Athens, Greece)

Friday 1:30 PM - 3:00 PM

Session 3: Open Discussion

Topic: Social Sites - Challenges and Opportunities

Moderator: Yannis Ioannidis (University of Athens, Greece)

Sihem Amer-Yahia (Qatar Computing Research Institute – QCRI, Qatar), Laks Lakshmanan (University of British Columbia, Canada), Amr El Abbadi (University of California Santa Barbara, USA), Raghu Ramakrishnan (Yahoo! Research, USA)

Friday 3:30 PM - 5:30 PM

Session 4: Personalization and Context-aware Models and Systems

- The CareDB Context and Preference-Aware Database System (Invited Paper) Justin Levandoski (Microsoft Research, USA), Mohamed Khalefa (Alexandria University, Egypt), Mohamed Mokbel (University of Minnesota, USA)
- Towards a Reference Model for Social User Profiles: Concept & Implementation Elisabeth Kapsammer (Johannes Kepler University Linz, Austria), Stefan Mitsch (Johannes Kepler University Linz, Austria), Birgit Pröll (Johannes Kepler University Linz, Austria), Werner Retschitzegger (Johannes Kepler University Linz, Austria), Wieland Schwinger (Johannes Kepler University Linz, Austria), Manuel Wimmer (Vienna University of Technology, Austria), Martin Wischenbart (Vienna University of Technology, Austria), Stefan Lechner (Netural Communication Linz, Austria)
- Problems and Opportunities in Context-based Personalization (Invited Paper) Letizia Tanca (Politecnico di Milano, Italy), Cristiana Bolchini (Politecnico di Milano, Italy), Elisa Quintarelli (Politecnico di Milano, Italy), Fabio A. Schreiber (Politecnico di Milano, Italy), Giorgio Orsi (University of Oxford, UK)
- Efficient Algorithms for Recommending Top-k Items and Packages (Invited Paper) Mohammad Khabbaz (University of British Columbia, Canada), Min Xie (University of British Columbia, Canada), Laks Lakshmanan (University of British Columbia, Canada)



VLDB 2011 Conference Officers

GENERAL CHAIR Per-Ake (Paul) Larson, Microsoft Research

PROGRAM CHAIRS

General PC chairs:

Jose Blakeley, Microsoft Joseph M. Hellerstein, University of California - Berkeley

Research track chairs:

Nick Koudas, University of Toronto Wolfgang Lehner, Dresden University of Technology Sunita Sarawagi, IIT Bombay

Industrial track chairs:

Berthold Reinwald, IBM Research - Almaden Phil Bohannon, Yahoo! Research

Challenges and Vision track chair: Gerhard Weikum, Max-Planck-Institut für Informatik

DEMO PROGRAM

Demo PC chairs: Jignesh Patel, University of Wisconsin, Madison Masatoshi Yoshikawa, University of Kyoto

TUTORIAL PROGRAM

Tutorial PC chairs:

Qiong Luo, Hong Kong University of Science & Technology Gerome Miklau, University of Massachusetts - Amherst

WORKSHOP PROGRAM

Workshop chairs: Luis Gravano, Columbia University Chris Jermaine, Rice University

PHD PROGRAM PhD Workshop chairs: Philippe Bonnet, IT University of Copenhagen

PROCEEDINGS

Proceedings chair: Uwe Roehm, University of Sydney

TEN-YEAR AWARDS COMMITTEE

Bettina Kemme, McGill University Tova Milo, Tel Aviv University Divesh Srivastava, AT&T Labs (Chair)

BEST PAPER AWARDS COMMITTEE

Umesh Dayal, HP Labs Donald Kossmann, ETH Zurich Stan Zdonik, Brown University (Chair)

SPONSORSHIPS

Dave Dewitt, Microsoft J. Gray Systems Lab

LOCAL ORGANIZING COMMITTEE

Chair: Christian König, Microsoft Research Treasurer: Magdalena Balazinska, University of Washington Web Master: Justin Levandoski, University of Minnesota Program Booklet Coordinator: Yasin Silva, Arizona State University Conference and Registration Management:

University of Washington Conference Management Team

VLDB 2011 participants will find that downtown Seattle is easily walkable and is waiting to be explored. A fine international restaurant community, a thriving retail shopping district, world-class cultural institutions, and many visitor attractions are within an easy walk of our meeting venue, the Westin Seattle Hotel. Seattle is surrounded by pristine waterways, the Olympic Mountains to the west, the Cascade Mountain range to the east and the majestic Mount Rainier to the south. To the east, a world-class wine region has grown up.

Seattle Highlights

Pike Place Farmers Market

The bounty of fresh produce and local artisan products at the friendly Pike Place Market has caused countless people to fall head over heels in love with Seattle. It's located right downtown and is a not to be missed. Pike Place's nine acres have been a staple in Seattle for more than a century. It's been called "the soul of Seattle," and for good reason. Be sure to see the original Starbucks store; the famous fish-throwers at Pike Place Fish Company; the colorful flower, produce and seafood stands on the main level; the kitschy and out-of-the ordinary shops on the lower level; and the row of heavenly scented food stands along Pike Place.

Seattle Waterfront

Take in the views of Elliott Bay and the Olympic Mountains from the Seattle waterfront. Take an inexpensive ferry ride to Bainbridge Island and enjoy the Seattle skyline from the bay. The waterfront offers both casual dining where you can grab some fish and chips or fine dining restaurants offering fresh, local seafood. Wander through the Olympic Sculpture Park on your way to the waterfront. It features 20-plus groundbreaking, large-scale sculptures in an awe-inspiring outdoor setting. It's free, by the way.

Seattle Central Library

This is the type of library that almost makes it hard to read — it's housed in a building that is just that striking. The design was conceived as five stacked boxes, which were staggered to allow maximum light into the building. Take photos inside the triangular Fifth Avenue entrance, which showcases a curtain of windows reflecting diamond-shaped light onto the floor.

[Designed by Rem Koolhaas and Joshua Ramus, Office for Metropolitan Architecture; built in 2004.]

Future of Flight Aviation Center and Boeing Tour Center

The Future of Flight Aviation Center and Boeing Tour Center is located in Mukilteo, Washington, 25 miles north of Seattle. Public tours of Boeing's Everett factory are available seven days a week. This facility is home to the 747, 767, 777 and 787 Dreamliner production lines. For reservations, call +1 360-756-0086 or toll free, +1 800-464-1476, from 9 a.m. to 5 p.m. (Pacific time). Ticket prices range from \$12 to \$20. www.boeing.com/commercial/tours/index.html

Seattle Center

The Seattle Center is home to many cultural and sports institutions, hosts annual festivals and features the iconic Space Needle. The grounds here were home to the 1962 World's Fair. Below are some of the highlights of the Seattle Center.

Seattle Center Monorail

What's a World's Fair without a little dash of the future? The monorail saw more than 7.3 million guests ride it during the fair's six month Seattle run. Now, the one-mile ride that takes two minutes from Westlake Center (at Fifth and Pine streets) to the Seattle Center carries about 2.5 million riders a year. The monorail departs every 10 minutes between 9 a.m. and 11 p.m. daily.

Space Needle

The Space Needle is two things: the most distinctive part of the Seattle skyline and the best place to view the city itself. With 360-degree views from the 605-foot-high observation deck, visitors can marvel at amazing views of Seattle and all the surrounding areas including the Puget Sound, the Olympic and Cascade mountains, Mount Rainier and more. There is also a gift shop at the bottom and SkyCity Restaurant right below the deck. Hours: Sunday through Thursday, 9 a.m. to 11 p.m. Friday and Saturday, 9 a.m. to midnight.

Experience Music Project Seattle Center

This museum is dedicated to rock music. So it's fitting that an early model took shape when Gehry deconstructed several electric guitars and used the pieces as building components. Take photos of the 21,000 metal shingles, each a unique shape, which appear to be different colors depending on how the light hits them. [Designed by Frank Gehry; built in 2000]

Pacific Science Center

With daily planetarium, science and laser shows, and two large IMAX theaters that showcase everything from Hollywood hits to educational features designed for the massive screens, the Pacific Science Center mixes fantastic entertainment with plenty of learning opportunities.

Bumbershoot (Annual Festival, September 3-5, 2011)

North America's largest urban arts festival, Bumbershoot[®]: Seattle's Music & Arts Festival takes place in the heart of the city at the 74-acre Seattle Center, Sept. 3-5. Since 1971, Bumbershoot has drawn artists representing the best in music, film, comedy, spoken word, dance, theatre, performance, and visual arts to the Seattle Center over Labor Day weekend. Over 100,000 visitors spend the weekend experiencing everything from regional favorites to international superstars across all arts disciplines. For more information, visit bumbershoot.org

Seattle Art Museum (SAM)

Wed–Sun: 10 am–5 pm; Thurs & Fri: 10 am–9 pm; Mon-Tues: closed. \$15 adult; \$12 senior (62+); \$9 student and teen (13-17); FREE: children 12 and under. Located downtown at 1300 First Avenue (1st Ave between Union and University Streets). Your admission includes access to visiting exhibitions as well as access to the SAM permanent collections and installations.

www.seattleartmuseum.com

Tours

On the Water

Taking one of the many water-based tours throughout Seattle will give a unique perspective on the city, and even the chance to peek at Bill Gates' massive lake estate. Outfits such as Argosy Cruises and Waterway Tours offers a variety of cruises including lunch and dinner trips on Elliott Bay, passage through the Chittenden Locks, cruises around the lakes, and themed trips. Another popular and unique option is Ride the Ducks of Seattle, which takes guests through city streets on an amphibious World War II vehicles, then into Lake Union for a look at beautiful houseboats, yachts and the skyline.

Seattle's Underground Tour

The city's history starts underground, where early developers struggled to build on the marshy land of what is now Pioneer Square. After numerous disasters and poor decisions, they decided to start over and build a new Seattle atop the old one. This whimsical tour takes guests through some of the areas left intact below the surface of today's streets. The day-time tours are g-rated for all crowds, but the adults-only nighttime Underworld Tour takes a decidedly seedier approach to the story telling.

Getting Around Seattle

King County Metro

Metro buses travel to nearly every corner of Seattle. They are easy to ride and offer an affordable and greener way to get around. For bus schedules and information, visit transit. metrokc.gov, which provides bus schedules and regional trip planning. Or call (206) 553-3000 24 hours a day.

The details of riding the bus:

- Exact fare is required at all times.
- Metro has a two-zone fare system. The Seattle city limits are the zone border.
- Generally, peak hours are Monday–Friday, 6–9 a.m. and 3–6 p.m.
- Nearly all Metro buses are fully accessible.
- Buses in the downtown zone are free 6 a.m.-7 p.m. daily.

The Transit Tunnel from Convention Place to the International District is open Monday–Saturday, 5–1 a.m. and Sunday, 6 a.m.-midnight.

Central Link Light Rail

New Link light rail trains now offer service from downtown Seattle to the Seattle-Tacoma International Airport, with eleven stops in between. Round-trip tickets from Sea-Tac International Airport to downtown's Westlake station cost \$5. Visit www.soundtransit.org for updates.

Streetcar

The South Lake Union Streetcar is the first in a planned streetcar network. It travels from downtown's Westlake Center to Fairview and Campus Dr. in the South Lake Union neighborhood, with 11 stops along the way. It runs seven days a week at 15-minute intervals during the following hours:

Monday–Thursday, 6 a.m.–9 p.m.; Friday–Saturday, 6 a.m.–11 p.m.; Sunday, 10 a.m.–7 p.m. Fares are \$2.50. For more information, go to www.seattlestreetcar.org.



© 2008 Seattle's Convention and Visitors Bureau



Runs **17x** Faster

1 Exadata ran

17x faster than

4 large UNIX servers.

ORACLE

oracle.com/goto/BNPparibas

BNP Paribas, 3rd largest global bank

Individual results depend on a number of factors. Actual results may vary. Copyright © 2010, Oracle and/or its affiliates. All rights reserved.

DATA CENTER TO **DEVICE** AND EVERYWHERE IN BETWEEN

It's your vital information and it's ready to work. In the office. On the road. At a customer site. Even a coffee shop. So manage it. Analyze it. Mobilize it. It's yours to use. Because in the Unwired Enterprise, your data is accessible anytime, anywhere.

Sybase is a proud sponsor of the 2011 VLDB Conference.

Copyright © 2011 Sybase, Inc. All rights reserved. Unpublished rights reserved under U.S. copyright laws. Sybase, and the Sybase logo, are trademarks of Sybase, Inc. or its subsidiares. "Indicates registration in the United States of America. SAP, and the SAP logo, are the trademarks or registered trademarks of SAP AG in Germany and in several other countries. All other trademarks are the property of their respective owners.





About SAP

As market leader in enterprise application software, SAP (NYSE: SAP) helps companies of all sizes and industries run better. From back office to boardroom, warehouse to storefront, desktop to mobile device – SAP empowers people and organizations to work together more efficiently and use business insight more effectively to stay ahead of the competition. SAP applications and services enable more than 170,000 customers (includes customers from the acquisition of Sybase) to operate profitably, adapt continuously, and grow sustainably. For more information, visit www.sap.com.

(*) SAP defines business software as comprising enterprise resource planning, business intelligence, and related applications.

About SAP Research

SAP Research is the global technology research unit of SAP. The group significantly contributes to SAP's product portfolio and extends its leading position in the market by identifying and shaping emerging IT trends and generating breakthrough technologies through applied research.

In contrast to SAP's product groups and development labs that work on new functions and releases, the researchers explore opportunities that haven't yet been developed into products.

SAP Research spreads its research and development activities around the globe. Currently, their thriving network comprises of 19 locations worldwide, including their headquarters in Walldorf, and numerous partners from the business and academic worlds. (Additional information is available at www.sap.com/research.)

Seattle - USA

SPONSOR - SILVER

DRIVING THE FUTURE OF BIG DATA ANALYTICS



Do you think BIG? Do you want to work on exciting and challenging problems? Do you want all that an R&D division has to offer: develop products that people use, build a patent portfolio, publish research? We are looking for you!

EMC Greenplum is committed to pioneering the future of big data analytics. We are hiring new and experienced Software Engineers at our HQ in Silicon Valley, CA, and our Centers of Excellence worldwide.

Check us out at: www.greenplum.com/about-us/jobs See who you know at Greenplum: www.linkedin.com/company/greenplum



IBM's rich history in information management encompasses revolutionary breakthroughs ranging

IBM Research

EMC²

from the invention of the relational database by researchers in the Silicon Valley to advanced text analytics capabilities demonstrated by Watson, which proved its computing prowess on the quiz show Jeopardy! earlier this year. Key technologies resulting from IBM research projects such as System R, R*, Garlic, DB2 pureXML, System S and Clio and System T, along with the invention of the first data mining algorithms, have shaped the industry.

IBM Research is the leading corporate research institution in the world, a global community of 3,000 world-class scientists who have been honored with five Nobel Prizes, six Turing Awards, five National Medals of Science, 80 memberships in the National Academies and 11 inductions into National Inventors Hall of Fame. For the eighteenth straight year, IBM received the most U.S. patents of any company worldwide.

IBM works with clients around the world to help their businesses grow "smarter" as the planet becomes more instrumented, interconnected and intelligent. Through its software, hardware, services, consulting and R&D expertise, IBM helps create systems that lead to less traffic, healthier food, cleaner water, and safer cities.



The Transaction Processing Performance Council (TPC) is a non-profit corporation founded to define transaction processing and database benchmarks and to disseminate objective, verifiable performance data to the industry. Over the past two decades, the TPC has been

integral in shaping the landscape of industry standard benchmarks and has had a significant impact on the computing industry's use of these standards.

The technology landscape is continually evolving and challenging industry experts and researchers to develop innovative techniques to evaluate and benchmark computing systems. The TPC remains committed to developing new benchmark standards and enhancing existing standards to keep pace. The sponsorship of the TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC) is directed towards achieving these goals.

The TPC is currently developing benchmark standards to address virtualization, data integration and data warehousing. Organizations that are interested in influencing the TPC benchmark development process are encouraged to become members. Additional information is available at http://www.tpc.org/





Silver Sponsor

Platinum Sponsor

