ACM SIGMOD / PODS / SOCC 2010

June 6th – 11th, 2010
Indianapolis, Indiana

http://www.sigmod2010.org/

TABLE OF CONTENTS

MESSAGE FROM THE GENERAL CHAIRS

It is our great pleasure to welcome you to the 2010 SIGMOD/PODS joint conference, which is being held in Indianapolis, the capital city of Indiana. Indianapolis is the nation's 13th largest city. It has gone through a dramatic revitalization and a stunning renaissance that makes it a different place than it was just a decade ago. Indy is the perfect balance of cosmopolitan style and small-town charm that makes it a successful destination for leisure travel, conventions and group tours, catering to nearly 22 million visitors a year. SIGMOD/PODS 2010 continues its tradition as a premier forum for the presentation of research results, industry developments, tutorials, and demos in the general area of data management.

Program chairs Divyakant Agrawal (SIGMOD), Dirk Van Gucht (PODS), Surajit Chaudhuri and Mendel Rosenblum (SOCC) with their program committees put together a rich technical program including keynotes, research papers, tutorials, demonstrations, industrial sessions, and panels. Overall, the program promises to be inspiring, educational, and exciting. The eight workshops with their varied topics as well as the Symposium on Cloud Computing make this year's meeting extra special.

Social events are an important part of any conference as a way to build friendships and collaborations. Our premier social function is always our banquet. This year the banquet is held in the Eiteljorg Museum of American Art. The Eiteljorg Museum is unique, one of only two museums east of the Mississippi with both Native American and Western art. Located in downtown Indianapolis, the museum is within walking distance of the conference hotel.

Many people worked hard on the organization of the joint conference. We would like to thank all members of our organizing committee for their time, efforts, and commitment to make this conference as good as it can be. We also would like to thank our industrial and academic sponsors for stepping up and continuing to provide strong support to our community even in the current economic downturn.

This year's conference is collocated with ACM Symposium on Cloud Computing 2010 (SOCC), the first in a new series of symposia co-sponsored by SIGMOD and SIGOPS, with the aim of bringing together researchers, developers, users, and practitioners interested in cloud computing. SOCC will be collocated with the SIGMOD and SOSP conferences in alternate years. There will be a reception for those registered for SOCC scheduled for Thursday evening at the Indiana Repertory Theater.

We hope that you will find our technical program interesting and thought-provoking, the social events entertaining, and that the conference will provide you with a valuable opportunity to share ideas with other researchers and practitioners from institutions around the world.

Sincerely,

Ahmed Elmagarmid (SIGMOD General Chair) , Jan Paredaens (PODS General Chair) & Joseph M. Hellerstein (SOCC General Chair)

SOCIAL EVENTS

PODS Reception
Sunday, June 6, 18:00-21:00
Hyatt Regency,
Cosmopolitan AB

Hyatt Regency Hotel

SOCIAL EVENTS

SIGMOD Reception
Monday, June 7, 18:30-21:00
Indiana Roof Ballroom

The Indiana roof ballroom, affectionately just called the "roof" has been an important part of downtown indianapolis for nearly a century. Located above the Indiana Repertory Theatre, it has hosted famous performers, high school proms, weddings, and ceremonies while creating the illusion of twilight in a Spanish town.

SOCIAL EVENTS:

Microsoft Reception
Tuesday, June 8, 20:30-23:00
Hyatt Regency Hotel

SOCIAL EVENTS:

Banquet
Wednesday, June 9, 18:30-23:30
Eiteljorg Museum of American Art

The Eiteljorg Museum, which opened in 1989, was founded by Harrison Eiteljorg. The museum showcases Western and Native American art and cultural objects. The museum's design is also inspired by the land, people and architecture of the American Southwest.
It is the only museum of its kind in the Midwest, and one of only two museums east of the Mississippi that showcase both Native American and Western art, culture and history.

.

SOCIAL EVENTS:

SOCC Reception
Thursday, June 10, 18:30-21:00
Indiana Repertory Theatre

INDIANA REPERTORY THEATRE

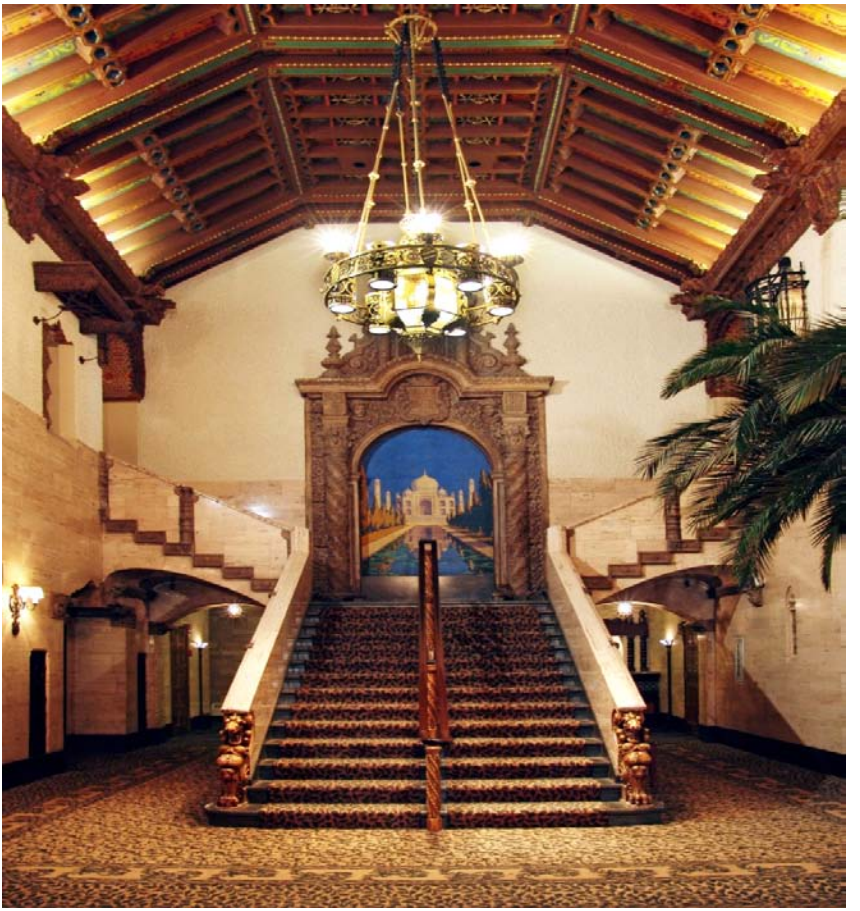

CONFERENCE AREA MAP

Transport Map

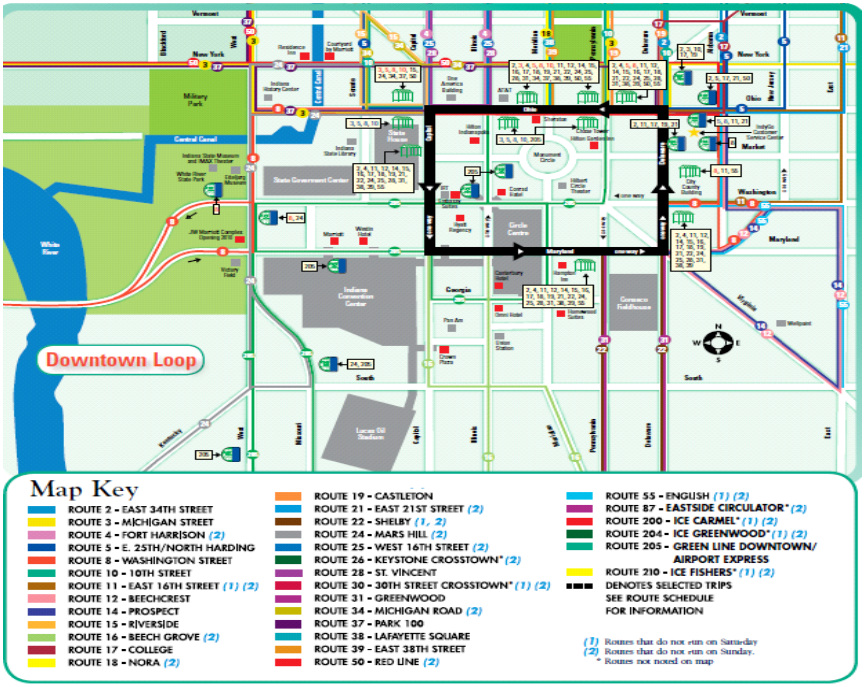

Hotel Layout

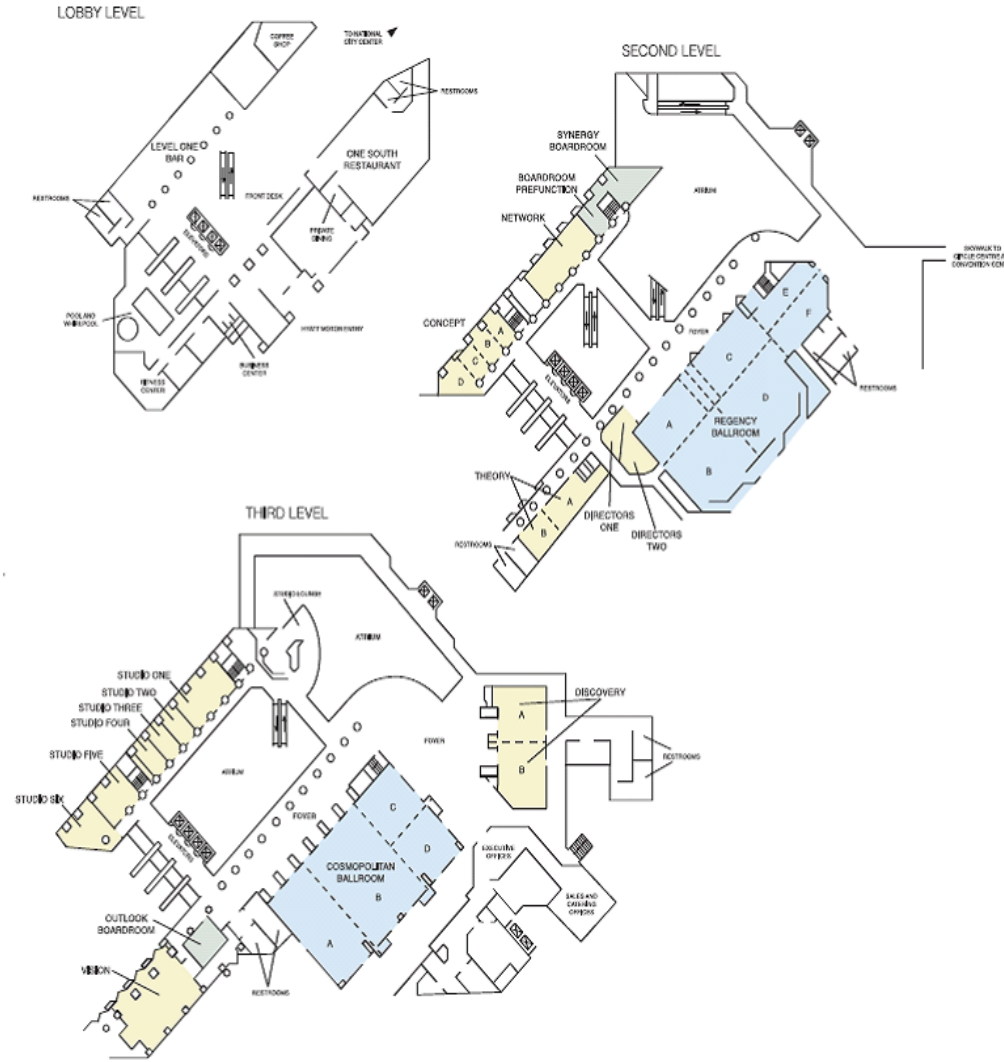

PROGRAM AT A GLANCE

SUNDAY AT A GLANCE				JUNE 6
8:00 – 8:30	Continental Breakfast (Cosmopolitan CD Foyer)			
8:30 – 10:00	WebDB (Cosmopolitan C)	KEYS (Cosmopolitan D)	MobiDE (Discovery A)	WANDS (Discovery B)
10:00 – 10:30	Coffee Break (Cosmopolitan CD Foyer)			
10:30 – 12:00	WebDB (Cosmopolitan C)	KEYS (Cosmopolitan D)	MobiDE (Discovery A)	WANDS (Discovery B)
12:00 – 13:30	Lunch Provided (Cosmopolitan AB)			
13:30 – 15:00	WebDB (Cosmopolitan C)	KEYS (Cosmopolitan D)	MobiDE (Discovery A)	WANDS (Discovery B)
15:00 – 15:30	Coffee Break (Cosmopolitan CD Foyer)			
15:30 – 17:00	WebDB (Cosmopolitan C)	KEYS (Cosmopolitan D)	MobiDE (Discovery A)	WANDS (Discovery B)
18:00 – 21:00	PODS Reception (Cosmpolitan AB)			

MONDAY AT A GLANCE			JUNE 7
8:00 – 8:30	Continental Breakfast (Regency A-D Foyer)		
8:30 – 9:45	PODS Session 1: Opening and Keynote (Ballroom A-D)		
9:45 – 10:15	Coffee Break (Regency A-D Foyer)		
10:15 – 11:45	PODS Session 2 (Regency Ballroom A-D)	DaMoN (Regency E)	DBTest (Regency F)
11:45 – 13:15	Lunch Provided for DaMoN and DBTest Participants (Cosmopolitan AB)		
13:15 – 14:30	PODS Session 2 (Regency Ballroom A-D)	DaMoN (Regency E)	DBTest (Regency F)
14:30 – 14:45	Break		
14:45 – 15:45	PODS Session 2 (Regency Ballroom A-D)	DaMoN (Regency E)	DBTest (Regency F)
15:45 – 16:15	Coffee Break (Regency A-D Foyer)		
16:15 – 18:15	PODS Session 2 (Regency Ballroom A-D)	DaMoN (Regency E)	DBTest (Regency F)
18:30 – 21:00	SIGMOD Reception and Undergrad Posters Competition (Indiana Roof Ballroom)		
21:00 – 23:00	PODS Business Meeting (Cosmopolitan CD)		

TUESDAY AT A GLANCE							JUNE 8
8:00 – 8:30	Continental Breakfast (Regency A-D Foyer)						
8:45 – 9:00	SIGMOD Welcome (Regency Ballroom A-D)						
9:00 – 10:00	SIGMOD Keynote 1 (Regency Ballroom A-D)						
10:00 – 10:30	Coffee Break (Regency and Cosmpolitan Foyer)						
10:30 – 12:00	PODS Session 6 (Regency EF)	SIGMOD Res 1 (Discovery)	SIGMOD Res 2 (Vision)	SIGMOD Res 3 (Cosmo-politan CD)	SIGMOD Tutorial 1 (Theory)	SIGMOD Industry 1 (Concept A-D)	SIGMOD Demo A (Studio One)
12:00 – 13:30	Lunch Buffet Provided (Cosmopolitan AB)						
13:30 – 15:00	PODS Session 7 (Regency EF)	SIGMOD Res 4 (Discovery)	SIGMOD Res 5 (Vision)	SIGMOD Res 6 (Cosmo-politan CD)	SIGMOD Tutorial 1 (Theory)	SIGMOD Industry 2 (Concept A-D)	SIGMOD Demo B & Prog. Contest (Studio One)
15:00 – 15:30	Coffee Break (Regency and Cosmpolitan Foyer)						
15:30 – 16:30	PODS Session 8 (Regency EF)	SIGMOD Res 7 (Discovery)	SIGMOD Res 8 (Vision)		SIGMOD Tutorial 2 (Theory)	Panel 1 (Regency Ballroom A-D)	SIGMOD Demo C (Studio One)
16:30 – 17:00	PODS Session 9 (Regency EF)						
17:00 – 17:30		Break					
18:00 – 20:30	New Researchers Symposium (Cosmopolitan CD)						
20:30 – 23:00	Microsoft Reception (Cosmopolitan AB)						

WEDNESDAY AT A GLANCE							JUNE 9
8:00 – 8:30	Continental Breakfast (Regency and Cosmpolitan Foyer)						
8:30 – 10:00	PODS Session 10 (Cosmo-politan A)	SIGMOD Res 9 (Discovery)	SIGMOD Res 10 (Vision)	SIGMOD Res 11 (Cosmo-politan CD)	SIGMOD Tutorial 3 (Theory)	SIGMOD Industry 3 (Concept A-D)	SIGMOD Demo D (Studio One)
10:00 – 10:30	Coffee Break (Regency and Cosmpolitan Foyer)						
10:30 – 12:00	PODS Session 11 (Cosmo-politan A)	SIGMOD Res 12 (Discovery)	SIGMOD Res 13 (Vision)	SIGMOD Res 14 (Cosmo-politan CD)	SIGMOD Tutorial 3 (Theory)	SIGMOD Industry 4 (Concept A-D)	SIGMOD Demo B & Prog. Contest (Studio One)
12:00 – 14:00	Provided Lunch and SIGMOD Business Meeting (Regency A-F)						
14:00 – 14:15	Break						
14:15 – 16:00	SIGMOD Awards Presentations (Regency A-F)						
16:00 – 17:30	PODS Session 12 (Cosmo-politan A)	SIGMOD Research Plenary I (Cosmopolitan BCD)					SIGMOD Demo A (Studio One)
17:30 – 18:00							Break
18:30 – 23:30	Banquet Eiteljorg Museum of American Art 18:30 hors d'oeuvres / 19:30 dinner						

THURSDAY AT A GLANCE						JUNE 10
8:30 – 9:00	Continental Breakfast (Regency A-D Foyer)					
9:00 – 10:00	SIGMOD Keynote Talk 2 (Regency Ballroom A-D)					
10:00 – 10:30	Coffee Break (Regency and Cosmpolitan Foyer)					
10:30 – 12:00	SIGMOD Res 15 (Discovery)	SIGMOD Res 16 (Vision)	SIGMOD Res 17 (Cosmo-politan CD)	SIGMOD Tutorial 4 (Theory)	SIGMOD Industry 5 (Concept A-D)	SIGMOD Demo C (Studio One)
12:00 – 13:30	Lunch on your own --- Oracle Dessert and Coffee Reception (Regency Ballroom AB)					
13:30 – 15:00	SIGMOD Res 18 (Discovery)	SIGMOD Res 19 (Vision)	SIGMOD Res 20 (Cosmo-politan CD)	SIGMOD Tutorial 4 (Theory)	SOCC (Regency EF)	SIGMOD Demo D (Studio One)
15:00 – 15:30	Coffee Break (Regency and Cosmpolitan Foyer)					
15:30 – 17:30	SIGMOD Research Plenary II (Cosmopolitan AB)			SOCC (Regency EF)		
17:30 – 17:45						
18:30 – 22:00	SOCC Reception Indiana Repertory Theater					

FRIDAY AT A GLANCE			JUNE 11
8:00 – 8:30	Continental Breakfast (Cosmopolitan CD Foyer)		
8:30 – 10:00	SOCC (Cosmopolitan CD)	Mentoring Wkshop (Discovery)	IDAR (Vision)
10:00 – 10:30	Coffee Break (Regency A-D Foyer)		
10:30 – 12:00	SOCC (Cosmopolitan CD)	Mentoring Wkshop (Discovery)	IDAR (Vision)
12:00 – 13:30	Coffee Break (Regency A-D Foyer)		
13:30 – 15:00	SOCC (Cosmopolitan CD)	Mentoring Wkshop (Discovery)	IDAR (Vision)
15:00 – 15:30	Coffee Break (Regency A-D Foyer)		
15:30 – 17:00	SOCC (Cosmopolitan CD)	Mentoring Wkshop (Discovery)	IDAR (Vision)

DETAILED CONFERENCE PROGRAM

SUNDAY 18:00 – 21:00

PODS Reception
Location: Cosmopolitan AB

MONDAY 8:30 – 9:45

PODS Opening and Keynote
Location: Ballroom A-D
Session Chair: Jan Paredaens (University of Antwerp)

Datalog Redux: Experience and Conjecture
Joseph M. Hellerstein (UC Berkeley)

MONDAY 10:15 – 11:45

PODS Session 2: Query Languages
Location: Ballroom A-D
Session Chair: Tova Milo (Tel Aviv University)

Expressive Languages for Path Queries over Graph-Structured Data
Pablo Barcela, Carlos Hurtado, Leonid Libkin and Peter Wood.
Transducing Markov Sequences
Benny Kimelfeld and Christopher Re
Positive Higher-Order Queries
Michael Benedikt, Gabriele Puppis and Huy Vu

MONDAY 13:15 – 14:30

PODS Session 3: Awards
Location: Ballroom A-D
Session Chair: Jianwen Su (UC Santa Barbara)

Alberto O. Mendelzon Test-of-Time Award: Typechecking for XML Transformers
Tova Milo, Dan Suciu, and Victor Vianu
Alberto O. Mendelzon Test-of-Time Award: Integrity Constraints for XML
Wenfei Fan and Jerome Simeon
Best Paper Award: An Optimal Algorithm for the Distinct Elements Problem
Daniel M. Kane, Jelani Nelson and David P. Woodruff
Regular Paper: Understanding Cardinality Estimation using Entropy Maximization
Christopher Re and Dan Suciu

MONDAY 14:45 – 15:45

PODS Session 4: Invited Tutorial 1
Location: Ballroom A-D
Chair: Marcelo Arenas (Pontifical Catholic Univ. of Chile)

From Information to Knowledge: Harvesting Entities and Relationships from Web Sources
Gerhard Weikum and Martin Theobald

MONDAY 16:15 – 18:15

PODS Session 5: Streams and Query Processing
Location: Ballroom A-D
Chair: Mikhail Atallah (Purdue University)

.

Optimal Sampling From Distributed Streams
Graham Cormode, Ke Yi, Qin Zhang and S. Muthukrishnan
Incremental Query Evaluation Revisited
Christoph Koch
Fast Manhattan Sketches in Data Streams
Jelani Nelson and David P. Woodruff
Semantic Query Optimization in the Presence of Types
Michael Meier, Michael Schmidt, Fang Wei and Georg Lausen

MONDAY 18:30 – 21:00

SIGMOD Reception and Undergrad Posters Competition
Location: Indiana Roof Ballroom

MONDAY 21:00 – 23:00

PODS Business Meeting
Location: Cosmopolitan CD

TUESDAY 8:45 – 10:00

SIGMOD Welcome and Keynote Talk 1
Location: Regency Ballroom A-D

The Flow of On-Line Information in Global Networks
Jon Kleinberg (Cornell Univesity)

TUESDAY 10:30 – 12:00

PODS Session 6: Privacy
Location: Regency EF
Chair: Graham Cormode (AT&T Labs Research)

Optimizing Linear Counting Queries under Differential Privacy
Chao Li, Michael Hay, Vibhor Rastogi, Gerome Miklau and Andrew McGregor
Universally Optimal Privacy Mechanisms for Minimax Agents
Mangesh Gupte and Mukund Sundararajan
Towards an Axiomatization of Statistical Privacy and Utility
Daniel Kifer and Bing-Rong Lin

SIGMOD Research Session 1: Advanced Query Processing
Location: Discovery
Chair: Walid Aref (Purdue University)

Efficiently Evaluating Complex Boolean Expressions
Marcus Fontoura, Yahoo! Research; Suhas Sadanandan, Yahoo! Inc; Jayavel Shanmugasundaram, Yahoo! Research; Sergei Vassilvitski, Yahoo! Research; Erik Vee, Yahoo! Research; Srihari Venkatesan, Yahoo! Inc; Jason Zien, Yahoo! Inc

How to ConQueR Why-Not Questions
Quoc Trung Tran, NUS; Chee-Yong Chan, National University of Singapore

Call to Order: A Hierarchical Browsing Approach to Eliciting Users' Preference
Feng Zhao, NUS; Gautam Das, University of Texas at Arlington; Kian-Lee Tan, National University of Singapore; Anthony Tung, National University of Singapore

Boosting Spatial Pruning: On Optimal Pruning of MBRs
Tobias Emrich, Ludwig-Maximilians-Universitaet Munchen; Hans-Peter Kriegel, University of Munich; Peer Kroger, Ludwig-Maximilians-Universitaet Munchen; Matthias Renz, Ludwig-Maximilians-Universitaet Munchen; Andreas Zuefle, Ludwig-Maximilians-Universitaet Munchen

SIGMOD Research Session 2: Databases on Modern Hardware
Location: Vision
Chair: Timos Sellis (National Technical University of Athens)

Leveraging Spatio-Temporal Redundancy for RFID Data Cleansing
Haiquan Chen, Auburn University; Wei-shinn Ku, Auburn University; Haixun Wang, Microsoft Research, Asia; Min-Te Sun, National Central University, Taiwan

Sampling Dirty Data for Matching Attributes
Henning Koehler, The University of Queensland; Xiaofang Zhou, The University of Queensland; Shazia Sadiq, The University of Queensland; Yanfeng Shu, CSIRO, Tasmanian ICT Centre; Kerry Taylor, CSIRO, ICT Centre

ERACER: A Database Approach for Statistical Inference and Data Cleaning
Chris Mayfield, Purdue University; Jennifer Neville, Purdue University; Sunil Prabhakar, Purdue University, USA

Recsplorer: Recommendation Algorithms based on Precedence Mining
Aditya Parameswaran, Stanford University; Georgia Koutrika, Stanford University; Benjamin Bercovitz, ; Hector Garcia-Molina, Stanford

SIGMOD Research Session 3: Graph Data and Querying
Location: Cosmopolitan CD
Chair: Lin Qiao, IBM Almaden Research

TEDI: Efficient Shortest Path Query Answering on Graphs
Fang Wei, University of Freiburg

GBLENDER: Towards Blending Visual Query Formulation and Query Processing in Graph Databases
Changjiu Jin, Nanyang Technological Univ; Sourav S Bhowmick, Nanyang Technological Univ; Xiaokui Xiao, NTU, Singapore; James Cheng, Nanyang Technological Univ; Byron Choi, Hong Kong Baptist University

Computing Label Constraint Reachability in Graph Databases
Ruoming Jin, Kent State University; Hui Hong, Kent State University; Haixun Wang, Microsoft Research, Asia; Yang Xiang, Kent State University; Ning Ruan, Kent State University

Pregel: A System for Large-Scale Graph Processing
Greg Malewicz, Google, Inc.; Matthew Austern, Google, Inc.; Aart Bik, Google, Inc.; James Dehnert, Google, Inc.; Ilan Horn, Google, Inc.; Naty Leiser, Google, Inc.; Grzegorz Czajkowski, Google, Inc.

SIGMOD Tutorial 1, Part I
Location: Theory

Mining Knowledge from Databases: An Information Network Analysis Approach
Jiawei Han, Yizhou Suny, Xifeng Yanz, and Philip S. Yu

SIGMOD Industrial Session 1: New Platforms
Location: Concept A-D
Chair: Divy Agarwal (UC Santa Barbara)

Experiences Evolving a New Analytical Platform: What Works and What's Missing
Jeff Hammerbacher (Cloudera)

Overview of SciDB: Large Scale Array Storage, Processing and Analysis
Paul Brown (SciDB)

Integrating Hadoop and parallel DBMS
Yu Xu (Teradata), Pekka Kostamaa (Teradata), Like Gao (Teradata)

A Comparison of Join Algorithms for Log Processing in MapReduce
Spyros Blanas (University of Wisconsin), Jignesh Patel (University of Wisconsin), Vuk Ercegovac, Jun Rao (IBM Research), Eugene Shekita (IBM Almaden Research Center), Yuanyuan Tian (IBM Almaden Research Center)

SIGMOD Demonstrations Session Group A
Location: Studio One

HadoopDB in Action: Building Real World Applications
Kamil Bajda-Pawlikowski, Yale University; Azza Abouzeid, Yale University; Jiewen Huang, Yale University; Daniel Abadi, Yale; Avi Silberschatz, Yale University

Online Aggregation and Continuous Query support in MapReduce
Tyson Condie, UC Berkeley; Neil Conway, UC Berkeley; Joseph Hellerstein, UC Berkeley; Peter Alvaro, UC Berkeley; John Gerth, Stanford University; Justin Talbot, Stanford University; Khaled Elmeleegy, Yahoo! Research; Russell Sears, Yahoo! Research

MapDupReducer: Detecting Near Duplicates over Massive Datasets
Chaokun Wang, Tsinghua University; Jianmin Wang, Tsinghua University China; Xuemin Lin, University of New South Wales; Wei Wang, University of New South Wales; Haixun Wang, Microsoft Research Asia; Hongsong Li, Microsoft Research Asia; Wanpeng Tian, Tsinghua University; Jun Xu, Tsinghua University; Rui Li, Tsinghua University

Large Graph Processing in the Cloud
Rishan Chen, MSRA; Xuetian Weng, MSRA; Bingsheng He, MSRA; Mao Yang, MSRA; Bo Peng, PKU

DCUBE: Discrimination Discovery in Databases
Salvatore Ruggieri, University of Pisa; Dino Pedreschi, Dipartimento di Informatica, Universita di Pisa; Franco Turini, Dipartimento di Informatica, Universita di Pisa

S-OLAP: an OLAP system for analyzing sequence data
Chun Kit Chui, The University of Hong Kong; Ben Kao, University of Hong Kong; Eric Lo, Hong Kong Polytechnic Universi; David Cheung, University of Hong Kong

ProgXe: Progressive Result Generation Framework for Multi-Criteria Decision Support Queries
Venkatesh Raghavan, Worcester Polytechnic Institut; Elke Rundensteiner, Worcester Polytechnic Institute

XTaGe: a flexible XML collection generator
Maria Perez, Universitat Jaume I; Ismael Sanz, Universitat Jaume I; Rafael Berlanga, Universitat Jaume I

K*SQL: A Unifying Engine for Sequence Patterns and XML
Barzan Mozafari, UCLA; Kai Zeng, UCLA; Carlo Zaniolo, UCLA

TUESDAY 13:30 – 15:00

PODS Session 7: Tutorial 2 and Information Complexity and Missing Information
Location: Regency EF

Chair: Jeffrey Naughton (Univ. of Wisconsin)

Tutorial 2: Information Complexity
T.S. Jayram

Regular Paper: Capturing Missing Tuples and Missing Values
Floris Geerts and Wenfei Fan.

SIGMOD Research Session 4: Data Streams and Time-series Data
Location: Discovery

Chair: Alex Labrinidis (University of Pittsburgh)

PR-Join: A Non-Blocking Join Achieving Higher Early Result Rate with Statistical Guarantees
Shimin Chen, Intel Labs Pittsburgh; Phillip Gibbons, Intel Labs Pittsburgh; Suman Nath, Microsoft

PODS: A New Model and Processing Algorithms for Uncertain Data Streams
Thanh Tran, UMass Amherst; Liping Peng, UMass Amherst; Boduo Li, UMass Amherst; Yanlei Diao, University of Massachusetts; Anna Liu, UMass Amherst

Fast Approximate Correlation for Massive Time-series Data
Abdullah Mueen, UC Riverside; Suman Nath, Microsoft; Jie Liu, Microsoft Research

An Algorithmic Approach to Event Summarization
Peng Wang, Fudan University; Haixun Wang, Microsoft; Majin Liu, Fudan University; Wei Wang, Fudan University

SIGMOD Research Session 5: Innovative Data Management
Location: Vision

Chair: Mirek Riedewald (Northeastern University)

Spreadsheet As a Relational Database Engine
Jerzy Tyszkiewicz, University of Warsaw

Scalable Architecture and Query Optimization for Transaction-time DBs with Evolving Schemas
Hyun Moon, NEC Labs; Carlo Curino, MIT; Carlo Zaniolo, UCLA

Data Conflict Resolution Using Trust Relationships
Wolfgang Gatterbauer, University of Washington; Dan Suciu, University of Washington

Analyzing the Energy Efficiency of a Database Server
Dimitris Tsirogiannis, University of Toronto; Stavros Harizopoulos, HP Labs; Mehul Shah, HP Labs

SIGMOD Research Session 6: Location and Sensor Based Data
Location: Cosmopolitan CD

Chair: Gottfried Vossen (WWU Munster)

Processing Proximity Relations in Road Networks
Zhengdao Xu, University of Toronto; Arno Jacobsen, University of Toronto

Searching Trajectories by Locations - An Efficiency Study
Zaiben Chen, The University of Queensland; Heng Tao Shen, University of Queensland, Australia; Xiaofang Zhou, The University of Queensland; Yu Zheng, Microsoft Research Asia; Xing Xie, Microsoft Research Asia

Processing Continuous Join Queries in Sensor Networks: a Filtering Approach
Mirco Stern, Universitaet Karlsruhe (TH); Erik Buchmann, Universitaet Karlsruhe (TH); Klemens Bohm, Universitaet Karlsruhe (TH)

TACO: Tunable Approximate Computation of Outliers in wireless sensor networks
Nikos Giatrakos, Unipi; Yannis Kotidis, Athens University of Economics and Business (AUEB) ; Antonios Deligiannakis, Technical University of Crete; Vasilis Vassalos, Athens University of Economics and Business; Yannis Theodoridis

SIGMOD Tutorial 1, Part II
Location: Theory

Mining Knowledge from Databases: An Information Network Analysis Approach
Jiawei Han, Yizhou Suny, Xifeng Yanz, and Philip S. Yu

SIGMOD Industrial Session 2: Exploiting New Hardware
Location: Concept A-D

Chair: Berthold Reinwald (IBM Almaden Research Center)

Ricardo: Integrating R and Hadoop
Yannis Sismanis (IBM Almaden), Sudipto Das (UC Santa Barbara), Rainer Gemulla (IBM Almaden Research Center), Peter Haas (IBM Almaden Research Center), Kevin Beyer (IBM Almaden Research Center), John McPherson (IBM Almaden Research Center)

PYMK: Friend Recommendation at MySpace
Michael Moricz (MySpace.com), Yerbolat Dosbayev (MySpace.com), Mikhail Berlyant (MySpace.com)

Forecasting High-Dimensional Data
Deepak Agarwal (Yahoo! Research), Datong Chen (Yahoo! Labs), Long-ji Lin (Yahoo! Labs), Jayavel Shanmugasundaram (Yahoo! Research), Erik Vee (Yahoo! Research)

Datawarehousing and Analytics Infrastructure at Facebook
Ashish Thusoo (Facebook), Dhruba Borthakur (Facebook)

SIGMOD Demonstrations Session Group B
Location: Studio One

Symbiote - A Reconfigurable Logic Assisted Data Stream Management System (RLADSMS)
Pranav Vaidya, IUPUI; Jaehwan John Lee, IUPUI; Fracis Bowen, IUPUI; Yingzi Du, IUPUI; Chandima Hewa Nadungodage, IUPUI; Yuni Xia, IUPUI

Interactive Visual Exploration of Neighbor-Based Patterns in Data Streams
Di Yang, WPI; Zhenyu Guo, WPI; Zaixian Xie, WPI; Elke Rundensteiner, Worcester Polytechnic Institute; Matthew Ward, WPI

TwitterMonitor: Trend Detection over the Twitter Stream
Michael Mathioudakis, University of Toronto; Nick Koudas, University of Toronto

Glacier: A Query-to-Hardware Compiler
Rene Mueller, ETH Zurich; Jens Teubner, ETH Zurich; Gustavo Alonso, ETH, Zurich

Exploratory Keyword Search on Data Graphs
Hilit Achiezra, The Hebrew University; Konstantin Golenberg, The Hebrew University; Benny Kimelfeld, IBM Almaden; Yehoshua Sagiv, The Hebrew University

Integrating Keyword Search with Multiple Dimension Tree Views over a Summary Corpus Data Cube
Mark Sifer, University of Wollongong; Yutaka Watanobe, University of Aizu, Japan; Subhash Bhalla, University of Aizu, Japan

Query Portals: Dynamically Generating Portals for Web Search Queries
Sanjay Agrawal, Microsoft; Kaushik Chakrabarti, MSR; Surajit Chaudhuri, Microsoft Research; Venkatesh Ganti, ; Christian Konig, Microsoft; Dong Xin, MSR

Creating and Exploring Web Form Repositories
Luciano Barbosa, University of Utah; Hoa Nguyen, University of Utah; Nguyen Thanh, University of Utah; Ramesh Pinnamaneni, University of Utah; Juliana Freire, University of Utah

TUESDAY 15:30 – 17:30

PODS Sessions 8 and 9: Uncertainty in Databases
Location: Regency EF

Chair: Benny Kimelfeld (IBM Almaden Research Center)

On the First-order Expressibility of Computing Certain Answers to Conjunctive Queries over Uncertain Databases
Jef Wijsen

Certain Answers for XML Queries
Claire David, Leonid Libkin and Filip Murlak
Computing Query Probability with Incidence Algebras
Nilesh Dalvi, Karl Schnaitter and Dan Suciu
On Probabilistic Fixpoint and Markov Chain Query Languages
Daniel Deutch, Cristoph Koch and Tova Milo

SIGMOD Research Session 7: Probabilistic and Uncertain Data
Location: Discovery, 15:30 – 17:00
Chair: Yannis Papakonstantinou (UC San Diego)

GRN Model of Probabilistic Databases: Construction, Transition and Querying
Ruiwen Chen, University of Ottawa; Yongyi Mao, University of Ottawa; Iluju Kiringa, University of Ottawa
Consistent Query Answers in Inconsistent Probabilistic Databases
Xiang Lian, HKUST; Lei Chen, Hong Kong University of Science and Technology; Shaoxu Song, HKUST
Threshold Query Optimization for Uncertain Data
Yinian Qi, Purdue University; Rohit Jain, Purdue University; Sarvjeet Singh, Purdue University; Sunil Prabhakar, Purdue University, USA;
Probabilistic String Similarity Joins
Jeffrey Jestes, Computer Science Department, FSU; Feifei Li, Florida State University; Zhepeng Yan, HKUST; Ke Yi, HKUST

SIGMOD Research Session 8: Leveraging Hardware for Data Management
Location: Vision, 15:30 – 17:00
Chair: Anastasia Ailamaki (EPFL)

FAST: Fast Architecture Sensitive Tree Search on Modern CPUs and GPUs
Changkyu Kim, Intel; Jatin Chhugani, Intel; Nadathur Satish, Intel Corporation; Eric Sedlar, Oracle; Anthony Nguyen, Intel; Tim Kaldewey, Oracle; Victor Lee, Intel Corporation; Scott Brandt, University of California, Santa Cruz; Pradeep Dubey, Intel
Fast In-Memory Sort on Modern CPUs and GPUs: A Case for Bandwidth-Oblivious SIMD Sort
Nadathur Satish, Intel Corporation; Changkyu Kim, Intel; Jatin Chhugani, Intel; Anthony Nguyen, Intel; Victor Lee, Intel Corporation; Daehyun Kim, Intel; Pradeep Dubey, Intel
Page-Differential Logging: An Efficient and DBMS-independent Approach for Storing Data into Flash Memory
Yi-Reun Kim, KAIST; Kyu-Young Whang, KAIST; Il-Yeol Song, Drexel University
Similarity Search and Locality Sensitive Hashing using Ternary Content Addressable Memories
Rajendra Shinde, Stanford University; Ashish Goel, Stanford University; Pankaj Gupta, Twitter Inc; Debojyoti Dutta, Cisco Systems Inc

SIGMOD Tutorial 2
Location: Theory, 15:30 – 17:00
Database Systems Research on Data Mining
Carlos Ordonez and Javier Garcia-Garcia

SIGMOD Panel 1: Crowds, Clouds, and Algorithms: Exploring the Human Side of "Big Data" Applications
Location: Regency Ballroom A-D, 15:30 – 17:00
Chair: Michael J. Franklin (UC Berkeley and Truviso, Inc.)
Panelists:
Sihem Amer-Yahia (Yahoo! Research),
AnHai Doan (Wisconsin),
Jon Kleinberg (Cornell),
Nick Koudas (U. Toronto and Sysomos, Inc.)

SIGMOD Demonstrations Session Group C
Location: Studio One, 15:30 – 17:00

Exploring Schema Similarity At Multiple Resolutions
Ken Smith, MITRE; Craig Bonaceto, MITRE; Chris Wolf, MITRE; Beth Yost, MITRE; Michael Morse, The MITRE Corporation; Peter Mork, The MITRE Corporation; Doug Burdick, MITRE

An Automated, yet Interactive and Portable DB designer
Ioannis Alagiannis, EPFL; Debabrata Dash, Carnegie Mellon University; Karl Schnaitter, UC Santa Cruz; Anastasia Ailamaki, EPFL; Neoklis Polyzotis, UC Santa Cruz

Midas: Integrating Public Financial Data
Sreeram Balakrishnan, IBM - Silicon Valley Lab.; Vivian Chu, IBM Research - Almaden; Mauricio Hernandez, IBM Reseach - Almaden; Howard Ho, IBM Research - Almaden; Rajasekar Krishnamurthy, ; Liu Shi, IBM Research - China; Jan Pieper, IBM Research - Almaden; Jeffrey Pierce, IBM Research - Almaden; Lucian Popa, IBM Research - Almaden; Christine Robson, IBM Research - Almaden; Lei Shi, IBM Research - China; Ioana Stanoi, IBM Research - Almaden; Edison Ting, IBM - Silicon Valley Lab.; Shivakumar Vaithyanathan, ; Huahai Yang, IBM Research - Almaden

Worry-Free Database Upgrades: Automated Model-Driven Evolution of Schemas and Complex Mappings
James Terwilliger, Microsoft Research; Philip Bernstein, Microsoft Corp.; Adi Unnithan, Microsoft Corporation

US-SQL: Managing Uncertain Schemata
Matteo Magnani, University of Bologna; Danilo Montesi, University of Bologna

PAROS: Pareto Optimal Route Selection
Matthias Schubert, Ludwig-Maximilians-University; Franz Graf, Ludwig-Maximilians-Universitaet Munchen; Matthias Renz, Ludwig-Maximilians-Universitaet Munchen; Hans-Peter Kriegel, University of Munich

MoveMine: Mining Moving Object Databases
Zhenhui Li, Univ. OF ILLINOIS AT URBANA-CH; Jae-Gil Lee, IBM Almaden Research Center; Jiawei Han, UIUC

PIQL: A Performance Insightful Query Language
Michael Armbrust, UC Berkeley; Stephen Tu, ; Armando Fox, UC Berkeley; Michael Franklin, UC Berkeley; David Patterson, UC Berkeley; Nick Lanham, UC Berkeley; Beth Trushkowsky, UC Berkeley; Jesse Trutna, UC Berkeley

DoCQS: A Prototype System for Supporting Data-oriented Content Query
Mianwei Zhou, UIUC; Tao Cheng, University of Illinois at Urbana-champaign; Kevin Chang, University of Illinois at Urbana-champaign

TUESDAY 18:00 – 20:30

SIGMOD New Researchers Symposium
Location: Cosmopolitan CD

TUESDAY 20:30 – 23:00

Microsoft Reception
Location: Cosmopolitan AB

WEDNESDAY 8:30 – 10:00

PODS Session 10: Schema Mappings and Design
Location: Cosmopolitan A
Chair: Jan Van den Bussche (Hasselt University)

Foundations of Schema Mapping Management
Marcelo Arenas, Jorge Pérez, Juan L. Reutter and Cristian Riveros

Schema Design for XML Repositories: Complexity and Tractability
Wim Martens, Matthias Niewerth and Thomas Schwentick

Simplifying XML Schema: Single-Type Approximations of Regular Tree Languages
Wouter Gelade, Tomasz Idziaszek, Wim Martens and Frank Neven

SIGMOD Research Session 9: Web Data Integration
Location: Discovery
Chair: Magdalena Balazinska (University of Washington)

Automatically Incorporating New Sources in Keyword Search-Based Data Integration
Partha Talukdar, University of Pennsylvania; Zachary Ives, University of Pennsylvania; Fernando Pereira, Google

Active Knowledge: Dynamically Enriching RDF Knowledge Bases by Web Services
Nicoleta Preda, Max-Planck Institute; Gjergji Kasneci, Microsoft Research; Fabian Suchanek, Microsoft Research; Thomas Neumann, Max-Planck Institute, Germany; Wenjun Yuan, Hong Kong University; Gerhard Weikum, Max-Planck Institute

Schema Clustering and Retrieval for Multi-domain Pay-As-You-Go Data Integration Systems
Hatem Mahmoud, University of Waterloo; Ashraf Aboulnaga, University of Waterloo

Expressive and Flexible Access to Web-Extracted Data: A Keyword-based Structured Query Language
Jeffrey Pound, University of Waterloo; Ihab Ilyas, U of Waterloo; Grant Weddell, University of Waterloo

SIGMOD Research Session 10: Social Networks and Community Data
Location: Vision
Chair: Susan Davidson (University of Pennsylvania)

Multiple Features Fusion for Social Media Applications
Bin Cui, Peking University; Anthony Tung, National University of Singapore; Ce Zhang, PKU; Zhe Zhao, PKU

Finding Maximal Cliques in Massive Networks by H*-Graph
James Cheng, ; Yiping Ke, CUHK; Ada Fu, CUHK; Jeffrey Xu Yu, Chinese University of Hong Kong; Linhong Zhu, NTU, Singapore

K-Isomorphism: Privacy Preservation in Network Publication against structural attack
James Cheng, ; Ada Fu, CUHK; Jia Liu, Chinese University of Hong Kong

Load-Balanced Query Dissemination in Democratic Communities
Emiran Curtmola, UCSD; Alin Deutsch, UCSD; K.K. Ramakrishnan, AT&T Research Labs; Divesh Srivastava, AT&T Labs - Research

SIGMOD Research Session 11: Scalable Data Analytics
Location: Cosmopolitan CD
Chair: Chris Olston (Yahoo! Research)

Automatic Contention Detection and Amelioration for Data-Intensive Operations
John Cieslewicz, Columbia University; Kenneth Ross, Columbia University; Kyoho Satsumi, Columbia University; Yang Ye, Columbia University

Efficient Parallel Set-Similarity Joins Using MapReduce
Rares Vernica, University of California, Irvine; Michael Carey, UC Irvine; Chen Li, Univ of California, Irvine and BiMaple

ParaTimer: A Progress Indicator for MapReduce DAGs
Kristi Morton, University of Washington; Magdalena Balazinska, University of Washington; Dan Grossman, University of Washington

The DataPath System: A Data-Centric Analytic Processing Engine for Large Data Warehouses
Subi Arumugam, U Florida; Alin Dobra, UFL; Christopher Jermaine, Rice U.; Luis Perez, Rice University; Niketan Pansare, Rice University

SIGMOD Tutorial 3, Part I
Location: Theory
Information theory for data management
Suresh Venkatasubramanian, and Divesh Srivastava

SIGMOD Industrial Session 3: Advances in DBMSs

Location: Concept A-D

Chair: Sunil Prabhakar (Purdue University)

Extreme Scale with Full SQL Language Support in Microsoft SQL Azure
Nigele Ellis (Microsoft), Gopal Kakivaya, Dave Campbell (Microsoft)

Pay-As-You-Go - an adaptive approach to provide full context aware text search over document content
Zhen Hua Liu (Oracle), Thomas Baby (Oracle), Sukhendu Chakraborty (Oracle), Junyan Ding (Oracle), Anguel Novoselsky (Oracle), Vikas Arora (Oracle)

Sedna: Native XML Database Management System (Internals Overview)
Dmitry Lizorkin (Institute for System Programming of the Russian Academy of Sciences)

Optimizing Tuple-store Query Execution
Scott Meyer (Metaweb Technologies Inc.), Jutta Degener (Metaweb Technologies Inc.), John Giannandrea (Metaweb Technologies Inc.), Barak Michener (Metaweb Technologies Inc.)

SIGMOD Demonstrations Session Group D

Location: Studio One

QRelX: Generating Meaningful Queries that Provide Cardinality Assurance
Manasi Vartak, Worcester Polytechnic Institute; Venkatesh Raghavan, Worcester Polytechnic Institut; Elke Rundensteiner, Worcester Polytechnic Institute

Performing Sound Flash Device Measurements: The uFLIP Experience
Matias Bjorling, University of Copenhagen; Lionel Le Folgoc, INRIA; Ahmed Mseddi, INRIA; Philippe Bonnet, University of Copenhagen; Luc Bouganim, INRIA; Bjorn Jonsson, Reykjavik University

GDR: A System for Guided Data Repair
Mohamed Yakout, Purdue University; Ahmed Elmagarmid, Purdue University; Jennifer Neville, Purdue University; Mourad Ouzzani, Purdue University

Crescando
Georgios Giannikis, ETH Zurich; Philipp Unterbrunner, ETH Zurich; Jeremy Meyer, Amadeus; Gustavo Alonso, "ETH, Zurich"; Dietmar Fauser, Amadeus; Donald Kossman, ETH Zurich

A Tool for configuring and visualizing database parameters
Vamsidhar Thummala, Duke Unviersity; Shivnath Babu, Duke University

Pluggable Personal Data Servers
Nicolas Anciaux, INRIA; Luc Bouganim, ; Yanli Guo, INRIA Paris-Rocquencourt; Philippe Pucheral, INRIA Paris-Rocquencourt; jean-jacques vandewalle, Gemalto; Shaoyi Yin, INRIA Paris-Rocquencourt

Mask: A System for Privacy-Preserving Policy-Based Access to Published Content
Mohamed Nabeel, Purdue University; Ning Shang, Purdue University; John Zage, Purdue University; Elisa Bertino,

SimDB: A Similarity-aware Database System
Yasin Silva, Purdue University; Walid Aref, Purdue University; Paul Larson, Microsoft Research

A Demonstration of FlexPref: Extensible Preference Evaluation Inside the DBMS Engine
Justin Levandoski, University of Minnesota; Mohamed F. Mokbel, Univ. of Minnesota; Mohamed Khalefa, University of Minnesota; Venkateshwar Korukanti, University of Minnesota

WEDNESDAY 10:30 – 12:00

PODS Session 11: Query Learning

Location: Cosmopolitan A

Chair: Dan Suciu (University of Washington)

Characterizing Schema Mappings via Data Examples
Bogdan Alexe, Phokion Kolaitis and Wang-Chiew Tan

Understanding Queries in a Search Database System
Ronald Fagin, Benny Kimelfeld, Yunyao Li, Sriram Raghavan and Shivakumar Vaithyanathan

A Learning Algorithm for Top-Down XML Transformations
Aurelien Lemay, Sebastian Maneth and Joachim Niehren

SIGMOD Research Session 12: Advanced Query Processing
Location: Discovery
Chair: Jiaheng Lu (Renmin University)

Variance Aware Optimization of Parameterized Queries
Surajit Chaudhuri, Microsoft Research; Hongrae Lee, University of British Columbia; Vivek Narasayya, Microsoft Research

Positional Update Handling in Column Stores
Sandor Heman, VectorWise; Marcin Zukowski, VectorWise; Niels Nes, ; Lefteris Sidirourgos, CWI; Peter Boncz, CWI

Durable Top-k Search in Document Archives
Leong Hou U, The University of Hong Kong; Nikos Mamoulis, University of Hong Kong; Klaus Berberich, MPII; Srikanta Bedathur, MPII

Ajax-based Report Pages as Incrementally Rendered Views
Yupeng FU, UCSD; Keith Kowalczykowski, app2you Inc; Yannis Papakonstantinou, UCSD; Kevin Keliang Zhao, UCSD; Kian Win Ong, UC San Diego

SIGMOD Research Session 13: Cloud Computing and Internet Scale Computing
Location: Vision
Chair: Mehul Shah (HP Labs)

An Evaluation of Alternative Architectures for Transaction Processing in the Cloud
Simon Loesing, ETH Zurich; Tim Kraska, ETH Zurich; Donald Kossmann, ETH Zurich

Indexing Multi-dimensional Data in a Cloud System
Jinbao Wang, Harbin Institute of Technology; Hong Gao, Harbin Institute of Technology; Sai Wu, National Univ. of Singapore; Beng chin Ooi, National University of Singapore

Low Overhead Concurrency Control in Partitioned DBMSs
Evan Jones, MIT; Daniel Abadi, Yale; Samuel Madden, MIT

Efficient Querying and Maintenance of Network Provenance at Internet-Scale
Wenchao Zhou, University of Pennsylvania; Micah Sherr, University of Pennsylvania; Tao Tao, University of Pennsylvania; Xiaozhou Li, University of Pennsylvania; Boon Thau Loo, University of Pennsylvania; Yun Mao, University of Pennsylvania

SIGMOD Research Session 14: Data Summarization
Location: Cosmopolitan CD
Chair: Lei Chen (Hong Kong University of Science & Technology)

Hierarchically Organized Skew-Tolerant Histograms for Geographic Data Objects
Yohan Roh, SAIT, Samsung Electronics; Jae Ho Kim, KAIST; Yon Dohn Chung, Korea University; Jin Hyun Son, Hanyang University; Myoung Ho Kim, KAIST

Logging Every Footstep: Quantile Summaries for the Entire History
Yufei Tao, Chinese University of Hong Kong; Ke Yi, HKUST; Sheng Cheng, CUHK; Jian Pei, Simon Fraser University; Feifei Li, Florida State University

Continuous Sampling for Online Aggregation Over Multiple Queries
Sai Wu, National Univ. of Singapore; Beng chin Ooi, National University of Singapore; Kian-Lee Tan, National University of Singapore

Histograms Reloaded: The Merits of Bucket Diversity
Carl-Christian Kanne, Univ. of Mannheim; Guido Moerkotte, University of Mannheim

SIGMOD Tutorial 3, Part II
Location: Theory

Information theory for data management
Suresh Venkatasubramanian, and Divesh Srivastava

SIGMOD Industrial Session 4: Information Integration
Location: Concept A-D
Chair: Chen Li (UC Irvine)

OpenII: An Open Source Information Integration Toolkit
Len Seligman (MITRE), Peter Mork (The MITRE Corporation), Alon Halevy (Google), Ken Smith (MITRE), Michael Carey (UC Irvine), Kuang Chen (University of California at Berkeley), Chris Wolf (MITRE), Jayant Madhavan (Google), Akshay Kannan (University of California at Berkeley)

Google Fusion Tables: Data Management, Integration and Collaboration in the Cloud
Jonathan Goldberg-Kidon (Google Inc.), Hector Gonzalez (Google Inc.), Alon Halevy (Google Inc.), Christian S. Jensen (Google Inc.), Anno Langen (Google Inc.), Jayant Madhavan (Google Inc.), Rebecca Shapely (Google Inc.)

Visual Interfaces to Data
Chris Stolte (Tableau Software)

Graphical XQuery in the AquaLogic Data Services Platform
Vinayak Borkar (University of California, Irvine), Michael Carey (UC Irvine), Sebu Koleth (Oracle), Alex Kotopoulis (Oracle), Kautul Mehta (SAP), Joshua Spiegel (Oracle), Sachin Thatte (Oracle), Till Westmann (SAP)

SIGMOD Demonstrations Session Group D and Programming Contest
Location: Studio One

For a list of Group D Demonstrations, please see Wednesday, 8:30-10:00.

WEDNESDAY 12:00 – 14:10

SIGMOD Business Meeting and Lunch
Location: Regency A-F

WEDNESDAY 14:15 – 16:00

SIGMOD Awards Presentation
Location: Regency A-F

WEDNESDAY 16:00 – 18:00

PODS Session 12: Constraints and Indexing
Location: Cosmopolitan A
Chair: Jan Hidders (Delft Univ. of Technology)

Cache-Oblivious Hashing
Rasmus Pagh, Zhewei Wei, Ke Yi and Qin Zhang.

Performance Guarantees for B-trees with Different-Sized Keys
Michael A. Bender, Haodong Hu and Bradley C. Kuszmaul.

When data dependencies over SQL tables meet the Logics of Paradox and S-3
Sven Hartmann and Sebastian Link.

The Power of Tree Projections: Local Consistency, Greedy Algorithms, and Larger Islands of Tractability
Gianluigi Greco and Francesco Scarcello

SIGMOD Research Plenary I
Location: Cosmopolitan BCD

Papers from SIGMOD Research Sessions 11-20.

SIGMOD Demonstrations Session Group A
Location: Studio One, 16:00 – 17:30

For a list of Group A Demonstrations, please see Tuesday, 10:30-12:00.

WEDNESDAY 18:30 – 23:30

Banquet at the Eiteljorg Museum of American Art

18:30 hors d'oeuvres / 19:30 dinner

THURSDAY 9:00 – 10:00

SIGMOD Keynote Talk 2
Location: Ballroom A-D
Chair: Anastasia Ailamaki (EPFL)

Warehouse Scale Computing
Luiz Barroso (Google)

THURSDAY 10:30 – 12:00

SIGMOD Research Session 15: Probabilistic Data, fuzzy Data, and Data Provenance
Location: Discovery
Chair: Martin Theobald (Max-Planck-Institut fur Informatik)

Lineage Processing over Correlated Probabilistic Databases
Bhargav Kanagal, University of Maryland; Amol Deshpande, Univ of Maryland
Monte Carlo Processing of Probabilistic Satisfiability Queries in MCDB
Luis Perez, Rice University; Subi Arumugam, U Florida; Christopher Jermaine, Rice U.
K-Nearest Neighbor Search for Fuzzy Objects
Kai Zheng, University of Queensland; Pui Cheong Fung, University of Queensland; Xiaofang Zhou, University of Queensland
An Optimal Labeling Scheme for Workflow Provenance Using Skeleton Labels
Zhuowei Bao, University of Pennsylvania; Susan Davidson, University of Pennsylvania; Sanjeev Khanna, University of Pennsylvania; Sudeepa Roy, University of Pennsylvania

SIGMOD Research Session 16: Data Security and Privacy
Location: Vision
Chair: Chris Clifton (Purdue University)

SecureBlox: Customizable Secure Distributed Data Processing
William Marczak, UC Berkeley; Shan Shan Huang, LogicBlox, Inc.; Martin Bravenboer, LogicBlox, Inc.; Micah Sherr, University of Pennsylvania; Boon Thau Loo, University of Pennsylvania; Molham Aref, LogicBlox
Differentially Private Aggregation of Distributed Time-Series with Transformation and Encryption
Vibhor Rastogi, University of Washington; Suman Nath, Microsoft
Non-homogeneous Generalization in Privacy Preserving Data Publishing
Wai Kit Wong, University of Hong Kong; Nikos Mamoulis, University of Hong Kong; David Cheung, University of Hong Kong
Preserving Privacy and Fairness in Peer-to-Peer Data Integration
Hazem Elmeleegy, Purdue University; Mourad Ouzzani, Purdue University; Ahmed Elmagarmid, Purdue University; Ahmad Abusalah, Purdue University

SIGMOD Research Session 17: Web Data Integration
Location: Cosmopolitan CD
Chair: Fatma Ozcan(IBM Almaden)

Structured Annotations of Web Queries
Nikos Sarkas, University of Toronto; Stelios Paparizos, Microsoft Research; Panayiotis Tsaparas, Microsoft Research

On Active Learning of Record Matching Packages
Arvind Arasu, Microsoft Research; Michaela Goetz, Cornell University; Raghav Kaushik, Microsoft Research

I4E: Interactive Investigation of Iterative Information Extraction
Anish Das Sarma, Yahoo Research; Alpa Jain, Yahoo; Divesh Srivastava, AT&T Labs - Research

ONDUX: On-Demand Unsupervised Learning for Information Extraction
Eli Vilarinho, Federal University of Amazonas; Altigran Silva, UFAM; Marcos Goncalves, UFMG; Edleno de Moura, Federal University of Amazonas

SIGMOD Tutorial 4, Part I
Location: Theory

Enterprise Information Extraction: Recent Developments and Open Challenges
Laura Chiticariu, Yunyao Li, Sriram Raghavan, and Frederick Reiss

SIGMOD Industrial Session 5: Transactions, Security, and Caching
Location: Concept A-D
Chair: Graham Cormode (AT&T Research)

Analytics over Continuous and DisContinuous (ACDC) Streams: The Truviso Approach
Sailesh Krishnamurthy (Truviso), Rushan Chen (Truviso), Jeffery Davis (Truviso), Daniel Farina (Truviso), Michael Franklin (Truviso), Alan Li (Truviso), Neil Thombre (Truviso)

IBM Infosphere Streams for Scalable, Real-time, Intelligent Transportation Services
Alain Biem (IBM TJ Watson), Eric Bouillet (IBM TJ Watson), Hanhua Feng (IBM TJ Watson), Anand Ranganathan (IBM TJ Watson), Anton Ribov (IBM TJ Watson), Olivier Verscheure (IBM TJ Watson), Haris Koutsopoulos (KTH), Carlos Moran (KTH)

SIE-OBI: A Streaming Information Extraction Platform for Operational Business Intelligence
Malu Castellanos (HP Labs), Chetan Gupta (HP Labs), Umesh Dayal (HP Labs), Song Wang (HP Labs)

SIGMOD Demonstrations Session Group C
Location: Studio One

For a list of Group C Demonstrations, please see Tuesday, 15:30-17:00.

THURSDAY 12:00 – 13:30

Lunch on your own --- Oracle Dessert and Coffee Reception

THURSDAY 13:30 – 15:00

SIGMOD Research Session 18: Web Data Management
Location: Discovery
Chair: Jun Tatemura (NEC Research Laboratories)

Optimizing Content Freshness of Relations Extracted From the Web Using Keyword Search
Mohan Yang, Shanghai Jiao Tong University; Haixun Wang, Microsoft Research, Asia; Lipyeow Lim, ; Min Wang, HP Labs

Feeding Frenzy: Selectively Materializing Users' Event Feeds
Adam Silberstein, Yahoo! Research; Jeffrey Terrace, Princeton University; Brian Cooper, Yahoo! Research; Raghu Ramakrishnan, Yahoo! Research

Constructing and Exploring Composite Items
Senjuti Basu Roy, UTA; Sihem Amer-Yahia, Yahoo! Research; Ashish Chawla, Yahoo! Inc; Gautam Das, University of Texas at Arlington; Cong Yu, Yahoo! Research

Unbiased estimation of size and other aggregates over hidden web databases
Arjun Dasgupta, University of Texas Arlington; Xin Jin, George Washington University; Bradley Jewell, University of Texas at Arlington; Nan Zhang, George Washington University; Gautam Das, University of Texas at Arlington

SIGMOD Research Session 19: Graph Mining
Location: Vision
Chair: Chen Li (UC Irvine)

Towards Proximity Pattern Mining in Large Graphs
Arijit Khan, ; Xifeng Yan, ; Kun-Lung Wu, IBM Watson Research Center

GAIA: Graph Classification Using Evolutionary Computation
Ning Jin, UNC at Chapel Hill; Calvin Young, UNC at Chapel Hill; Wei Wang, University of North Carolina at Chapel Hill

Finding Maximum Degrees in Hidden Bipartite Graphs
Yufei Tao, Chinese University of Hong Kong; Sheng Cheng, CUHK; Jianzhong Li, Harbin Institute of Technology

Connected Substructure Similarity Search
Haichuan Shang, UNSW; Xuemin Lin, University of New South Wales; Wei Wang, University of New South Wales; Jeffrey Xu Yu, Chinese University of Hong Kong; Ying Zhang, UNSW

SIGMOD Research Session 20: Indexing and Storage Management
Location: Cosmopolitan CD
Chair: Daniel Abadi (Yale University)

B^ed-Tree: An All-Purpose Tree Index for String Similarity Search on Edit Distance
Zhenjie Zhang, National University of Singapo; Beng chin Ooi, National University of Singapore; Marios Hadjieleftheriou, AT&T Labs - Research; Divesh Srivastava, AT&T Labs - Research

On Indexing Error-Tolerant Set Containment
Raghav Kaushik, Microsoft Research; Parag Agrawal, Stanford University; Arvind Arasu, Microsoft Research

Workload-Aware Storage Layout for Database Systems
Oguzhan Ozmen, University of Waterloo; Kenneth Salem, University of Waterloo; Jiri Schindler, NetApp, Inc.; Steve Daniel, NetApp, Inc.

Querying Data Provenance
Grigoris Karvounarakis, University of Pennsylvania; Zachary Ives, University of Pennsylvania, Val Tannen, University of Pennsylvania

SIGMOD Tutorial 4, Part II
Location: Theory

Enterprise Information Extraction: Recent Developments and Open Challenges
Laura Chiticariu, Yunyao Li, Sriram Raghavan, and Frederick Reiss

SOCC Welcome
Location:Regency EF, 13:30 – 13:40

SOCC Keynote 1
Location:Regency EF, 13:40 – 14:40
Chair: Mendel Rosenblum (Stanford University)

Evolution and Future Directions of Large-Scale Storage and Computation Systems at Google
Jeffrey Dean (Google)

SOCC Session 1: Operating Systems
Location:Regency EF, 14:40 – 15:00

Chair: Mendel Rosenblum (Stanford University)

An Operating System for Multicore and Clouds: Mechanisms and Implementation
David Wentzlaff (MIT) , Charles Gruenwald III (MIT CSAIL) , Nathan Beckmann (MIT CSAIL) , Kevin Modzelewski (MIT CSAIL) , Adam Belay (MIT CSAIL) , Lamia Youseff (MIT CSAIL) , Jason Miller (MIT CSAIL) , Anant Agarwal (MIT CSAIL)

THURSDAY 15:30 – 17:45

SIGMOD Research Plenary I
Location: Cosmopolitan AB, 15:30 – 17:30

Papers from SIGMOD Research Sessions 1-10.

SOCC Session 2: Virtualization
Location:Regency EF, 15:30 – 16:30
Session Chair: Jeff Hammerbacher (Cloudera)

Lithium: Virtual Machine Storage for the Cloud
Jacob Hansen (VMware) , Eric Jul (Bell Labs, Dublin)
Differential Virtual Time (DVT): Rethinking I/O Service Differentiation for Virtual Machines
Mukil Kesavan (Georgia Institute of Technology), Ada Gavrilovska (Georgia Institute of Technology), Karsten Schwan (Georgia Institute of Technology)
Virtual Machine Power Metering and Provisioning
Aman Kansal (Microsoft Research) , Feng Zhao (Microsoft Research) , Jie Liu (Microsoft Research), Nupur Kothari (USC) , Arka Bhattacharya (IIT Kharagpur)

SOCC Session 3: Distributed and Parallel Processing
Location:Regency EF, 16:30 – 17:30
Session Chair: Brian Bershad (Google)

Stateful Bulk Processing for Incremental Algorithms
Dionysios Logothetis (UC San Diego), Christopher Olston (Yahoo! Research) , Benjamin Reed (Yahoo! Research) , Kevin Webb (UC San Diego) , Kenneth Yocum (UC San Diego)
Comet: Batched Stream Processing for Data Intensive Distributed Computing
Bingsheng He (Microsoft Research), Mao Yang (Microsoft Research) , Zhenyu Guo (Microsoft Research) , Rishan Chen (Beijing University) , Wei Lin (Microsoft Research) , Bing Su (Microsoft Research) , lidong Zhou (Microsoft Research)
Skew-Resistant Parallel Processing of Feature-Extracting Scientific User-Defined Functions
YongChul Kwon (University of Washington) , Balazinska Magdalena (University of Washington) , Bill Howe (University of Washington) , Jerome Rolia (HP)

THURSDAY 18:30 – 22:00

SOCC Reception at the Indiana Repertory Theater

FRIDAY 8:30 – 17:30

SOCC Keynote 2
Location:Cosmopolitan CD, 8:30 – 9:30
Session Chair: Surajit Chaudhuri (Microsoft Research)

Building Facebook: Performance at Massive Scale
Jason Sobel (Facebook)

SOCC Session 4: Applications
Location:Cosmopolitan CD, 9:30 – 10:00
Session Chair: Surajit Chaudhuri (Microsoft Research)

Hermes: Clustering Users in Large-Scale E-mail Services
Thomas Karagiannis (Microsoft Research), Christos Gkantsidis (Microsoft Research), Dushyanth Narayanan (Microsoft Research) , Antony Rowstron (Microsoft Research)

Defining Future Platform Requirements for e-Science Clouds
Lavanya Ramakrishnan (Lawrence Berkeley National Lab), Keith Jackson (Lawrence Berkeley National Lab), Shane Canon (Lawrence Berkeley National Lab), Shreyas Cholia (Laurence Berkeley National Lab), John Shalf (Lawrence Berkeley National Lab) [Position paper]

SOCC Session 5: Programming Models and Optimization
Location:Cosmopolitan CD, 10:30 – 11:30
Session Chair: Armando Fox (UC Berkeley)

Fluxo: A System for Internet Service Programming by Non-expert Developers
Emre Kiciman (Microsoft Research), Benjamin Livshits (Microsoft Research), Madanlal Musuvathi (Microsoft Research)

Nephele/PACs: A Programming Model and Execution Framework for Web-Scale Analytical Processing
Stephan Ewen (TU Berlin) , Fabian Hueske (TU Berlin) , Daniel Warneke (TU Berlin) , Dominic Battré (TU Berlin) , Volker Markl (TU Berlin) , Odej Kao (TU Berlin)

The Case For PIQL: A Performance Insightful Query Language
Michael Armbrust (UC Berkeley), Nick Lanham (UC Berkeley), Stephen Tu (UC Berkeley) , Armando Fox (UC Berkeley) , Michael Franklin (UC Berkeley) , David Patterson (UC Berkeley) [Position paper]

Towards Automatic Optimization of MapReduce Programs
Shivnath Babu (Duke University) [Position Paper]

SOCC Session 6: Benchmarking and Testing
Location:Cosmopolitan CD, 11:30 – 12:00
Session Chair: Armando Fox (UC Berkeley)

Benchmarking Cloud Serving Systems with YCSB
Brian Cooper (Yahoo! Research) , Adam Silberstein (Yahoo! Research) , Erwin Tam (Yahoo! Research) , Raghu Ramakrishnan (Yahoo!) , Russell Sears (Yahoo! Research)

Automated Software Testing as a Service
George Candea (EPFL) , Stefan Bucur (EPFL) , Cristian Zamfir (EPFL) [Position Paper]

SOCC Keynote 3

Location:Cosmopolitan CD, 13:30 – 14:30

Session Chair: Joseph M. Hellerstein (UC Berkeley)

The Internal Design of Salesforce.com's Multi-Tenant Architecture
Rob Woollen (Salesforce.com)

SOCC Session 7: Data Services

Location:Cosmopolitan CD, 14:30 – 15:00

Session Chair: Joseph M. Hellerstein (UC Berkeley)

G-Store: A Scalable Data Store for Transactional Multi key Access in the Cloud
Sudipto Das (UC Santa Barbara) , Divyakant Agrawal (UC Santa Barbara) , Amr El Abbadi (UC Santa Barbara)

Google Fusion Tables: Data Management, Integration and Collaboration in the Cloud
Alon Halevy (Google) , Hector Gonzalez (Google) , Jayant Madhavan (Google) , Christian S. Jensen (Aalborg University) , Jonathan Goldberg-Kidon (MIT) , Warren Shen (Google) , Rebecca Shapley (Google) , Anno Langen (Google) [Industrial Presentation]

SOCC Session 8: High Availability and Reliability

Location:Cosmopolitan CD, 15:30 – 16:30

Session Chair: Brian F. Cooper (Yahoo!)

Making Cloud Intermediate Data Fault-Tolerant
Steve Ko (Princeton University) , Imranul Hoque (University of Illinois at Urbana-Champaign) , Brian Cho (University of Illinois at Urbana-Champaign), Indranil Gupta (University of Illinois at Urbana-Champaign)

Characterizing Cloud Computing Hardware Reliability
Kashi Vishwanath (Microsoft Research) , Nachi Nagappan (Microsoft Research)

A Self-Organized, Fault-Tolerant and Scalable Replication scheme for Cloud Storage
Nicolas Bonvin (EPFL), Thanasis Papaioannou (EPFL), Karl Aberer (EPFL)

SOCC Session 9: Storage and System Modelling

Location:Cosmopolitan CD, 16:30 – 17:30

Session Chair: Michael Carey (UC Irvine)

Robust and Flexible Power-Proportional Storage
Hrishikesh Amur (Georgia Institute of Technology), James Cipar (Carnegie Mellon University) , Varun Gupta (Carnegie Mellon University) , Michael Kozuch (Intel Corporation) , Gregory Ganger (Carnegie Mellon University) , Karsten Schwan (Georgia Institute of Technology)

RACS: A Case for Cloud Storage Diversity
Lonnie Princehouse (Cornell University) , Hussam Abu-Libdeh (Cornell University) , Hakim Weatherspoon (Cornell University)

Characterizing, Modeling, and Generating Workload Spikes for Stateful Services
Peter Bodik (UC Berkeley) , Armando Fox (UC Berkeley) , Michael Franklin (UC Berkeley) , Michael Jordan (UC Berkeley) , David Patterson (UC Berkeley)

INVITED TALKS AND KEYNOTES

PODS Opening and Keynote: Datalog Redux: Experience and Conjecture

Joseph M. Hellerstein, UC Berkeley

Monday, 8:30 – 9:45

Location: Regency Ballroom A-D

Datalog was a foundational topic in the early years of PODS, despite skepticism from practitioners about its relevance. This has been changing in recent years, with unlikely champions exploring and promoting its use as a practical basis for programming in a wide variety of application domains. We reflect on our use of Datalog to build systems of significant complexity for both networking and cloud computing infrastructure. Based on that experience, we present conjectures regarding next-generation programming languages, and the role that database theory could play in their development.

Joseph M. Hellerstein is a Professor of Computer Science at the University of California, Berkeley, whose research focuses on data management and distributed systems. He is an ACM Fellow, and his work has been recognized by multiple awards including an Alfred P. Sloan Research Fellowship, and two ACM-SIGMOD "Test of Time" awards. Key ideas from his research have been incorporated into commercial and open-source database software released by IBM, Oracle, and PostgreSQL. He has also held industrial posts including Director of Intel Research Berkeley and Chief Scientist of Cohera Corporation, and currently serves as an advisor to a number of technology companies.

PODS Invited Tutorial 1: From Information to Knowledge: Harvesting Entities and Relationships from Web Sources

Gerhard Weikum and Martin Theobald, Max Planck Institute for Informatics, Saarbruecken, Germany

Monday, 14:45 – 15:45

Location: Regency Ballroom A-D

There are major trends to advance the functionality of search engines to a more expressive semantic level. This is enabled by the advent of knowledge-sharing communities such as Wikipedia and the progress in automatically extracting entities and relationships from semistructured as well as natural-language Web sources. Recent endeavors of this kind include DBpedia, EntityCube, KnowItAll, ReadTheWeb, and our own YAGO-NAGA project (and others). The goal is to automatically construct and maintain a comprehensive knowledge base of facts about named entities, their semantic classes, and their mutual relations as well as temporal contexts, with high precision and high recall. This tutorial discusses state-of-the-art methods, research opportunities, and open challenges along this avenue of knowledge harvesting. It also addresses issues of querying the knowledge base and ranking answers.

Gerhard Weikum is a Scientific Director at the Max-Planck Institute for Informatics, where he is leading the research group on databases and information systems. Earlier he held positions at Saarland University in Germany, ETH Zurich in Switzerland, MCC in Austin, and he was a visiting senior researcher at Microsoft Research in Redmond. His recent working areas include peer-to-peer information systems, the integration of database-systems and information-retrieval methods, and information extraction for building and maintaining large-scale knowledge bases. Gerhard has co-authored more than 150 publications, including a comprehensive textbook on transactional concurrency control and

recovery. He received the VLDB 2002 ten-year award for his work on self-tuning databases, and he is an ACM Fellow. He is a member of the German Academy of Science and Engineering and a member of the German Council of Science and Humanities. Gerhard has served on the editorial boards of various journals including ACM TODS and the new CACM, and as program committee chair for conferences like ICDE 2000, SIGMOD 2004, CIDR 2007, and ICDE 2010. From 2004 to 2009 he was president of the VLDB Endowment.

Martin Theobald is a Senior Researcher at the Max Planck Institute for Informatics. He obtained a doctoral degree in computer science from Saarland University, and spent two years as a post-doc at Stanford University where he worked on the Trio probabilistic database system. Martin received an ACM SIGMOD dissertation award honorable mention in 2006 for his work on the TopX search engine for efficient ranked retrieval of semistructured XML data.

PODS Invited Tutorial 2: Information Complexity

T. S. Jayram, IBM Almaden Research Center

Tuesday, 13:30 – 14:30

Location: Regency EF

The recent years have witnessed the overwhelming success of algorithms that operate on massive data. Several computing paradigms have been proposed for massive data set algorithms such as data streams, sketching, sampling etc. and understanding their limitations is a fundamental theoretical challenge. In this survey, we describe the information complexity paradigm that has proved successful in obtaining tight lower bounds for several well-known problems. Information complexity quantifies the amount of information about the inputs that must be necessarily propagated by any algorithm in solving a problem. We describe the key ideas of this paradigm, and highlight the beautiful interplay of techniques arising from diverse areas such as information theory, statistics and geometry.

T.S. Jayram is a manager of the Algorithms and Computation group at IBM Almaden Research Center. His research interests are in massive data sets, probabilistic databases and computational complexity. He obtained his Ph.D. from University of Washington in 1998 and has been with IBM Research since that time. He has received the IBM Outstanding Innovation Award for his contributions to massive data sets. Previously, he received the Machtey award for best student paper in FOCS 1995.

SIGMOD Keynote Talk 1: The Flow of On-Line Information in Global Networks

Jon Kleinberg, Cornell University

Tuesday, 9:00 – 10:00

Location: Regency Ballroom A-D

Increasingly, the information we experience on-line comes to us continuously over time, assembled from many small pieces, and conveyed through our social networks. This merging of information, network structure, and flow over time requires new ways of reasoning about the large-scale behavior of global networks, drawing on analogies to notions including contagion and mutation.

We discuss a set of approaches for tracking information as it travels and mutates in on-line networks, applying these ideas to a set of related problems. First, we show how this type of analysis can capture temporal patterns in the news over a daily time-scale --- in particular, the succession of story lines that evolve, compete for attention, and collectively produce an effect that commentators refer to as the `news cycle.' Second, we show how this approach can be combined with an analysis of network structure to trace the diffusion of specific pieces of information as they spread between people at a global scale.

This talk includes joint work with Lars Backstrom, Jure Leskovec, and David Liben-Nowell.

Jon Kleinberg is on the faculty of the Computer Science Department at Cornell University, where he holds the position of Tisch University Professor. His research focuses on issues at the interface of networks and information, with an emphasis on the social and information networks that underpin the Web and other on-line media. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences, and the recipient of MacArthur, Packard, and Sloan Foundation Fellowships, the Nevanlinna Prize, Katayanagi Prize, ACM-Infosys Foundation Award, and National Academy of Sciences Award for Initiatives in Research.

SIGMOD Keynote Talk 2: Warehouse Scale Computing

Luis Barroso, Google

Thursday, 9:00 – 10:00

Location: Regency Ballroom A-D

Outsourcing data storage and processing to large datacenter facilities is an attractive model for offering computing services, due in part to the efficiencies achievable by co-location of vast computing and storage capabilities and by amortizing their cost over many users and applications. Realizing those efficiencies, however, requires tackling hardware and software design challenges of an unprecedented scale. In this talk I will go over a small number of such challenges, with a focus on energy efficiency and distributed storage systems.

Luiz Andre Barroso is a Distinguished Engineer at Google, where he has worked across several engineering areas, ranging from applications and software infrastructure to the design of Google's computing platform. He was previously a member of the research staff at Digital Equipment Corporation (later Compaq), where his group did some of the early work on chip-multiprocessing systems including the design of Piranha, an architecture that helped inspire some of today'smulti-core CPUs. Luiz has a Ph.D. degree in computer engineering from the University of Southern California and B.S. and M.S. degrees in electrical engineering from the Pontifícia Universidade Católica, Rio de Janeiro.

SOCC Keynote Talk 1: Evolution and Future Directions of Large-Scale Storage and Computation Systems at Google

Jeffrey Dean, Google

Thursday, 1:40 – 2:40

Location: Regency Ballroom A-D

Underlying the many products and services offered by Google is a collection of systems and tools that simplify the storage and processing of large-scale data sets. These systems are intended to work well in Google's computational environment of large numbers of commodity machines connected by commodity networking hardware. Our systems handle issues like storage reliability and availability in the face of machine failures, and our processing tools make it relatively easy to write robust computations that run reliably and efficiently on thousands of machines. In this talk I'll highlight some of the systems we have built and are currently developing, and discuss some challenges and future directions for new systems.

Jeff joined Google in 1999 and is currently a Google Fellow in Google's Systems Infrastructure Group. He has co-designed/implemented five generations of Google's crawling, indexing, and query serving systems, and co-designed/implemented major pieces of Google's initial advertising and AdSense for Content systems. He is also a co-designer and co-implementor of Google's distributed computing infrastructure, including the MapReduce and BigTable systems, has worked on system software for statistical machine translation, and implemented a variety of internal and external developer tools.

SOCC Keynote Talk 2: Building Facebook: Performance at Massive Scale

Jason Sobel, Facebook

Friday, 8:30 – 9:30

Location: Cosmopolitan CD

From the day that Mark Zuckerberg started building Facebook in his Harvard dorm room in 2004 to today, the site has been built on common open source software such as Linux, Apache, MySQL, and PHP. Today Facebook reaches over 400 million people per month, is the largest PHP site in the World, and has released major pieces of our infrastructure as open source. It's not possible to scale a site like Facebook simply by sharding your databases, rather we've developed and contributed to a series of open source infrastructure technologies. Some of these projects include Cassandra, Hive, Haystack, memcached, and Scribe, where each focuses on solving a specific problem with Thrift allowing them to communicate across languages. This talk will give you a better idea of what it takes to scale Facebook, a look into the infrastructure we use to do so, and dive into our data model and new systems we've built for querying that data.

Jason Sobel is an engineering manager on the Infrastructure team at Facebook where he focuses on keeping the site fast and scalable. He helped develop Facebook's multiple datacenter strategy and also worked on Haystack, Facebook's custom built photo storage infrastructure. Before joining Facebook Jason was a member of the WAFL team at NetApp where he worked on filesystem performance. He holds bachelor's and master's degrees in computer science from Brown University.

SOCC Keynote Talk 3: The Internal Design of Salesforce.com's Multi-Tenant Architecture

Rob Woollen, Salesforce.com

Friday, 1:30 – 2:30

Location: Cosmopolitan CD

Salesforce.com provides a scalable multi-tenant cloud platform for more than 135,000 applications. This talk details the internal architecture behind the cloud including our shared multi-tenant resources, query processing and optimization, metadata-driven applications, and procedural development via Apex code.

Rob Woollen is a Principal Architect at Salesforce.com. At Salesforce Rob leads the architecture for Salesforce Chatter, a real-time collaboration platform. Prior to Salesforce, Rob spent 9 years with BEA Systems serving as the WebLogic Server Chief Architect. Rob holds a BSE in computer science from Princeton University.

AWARDS

SIGMOD Edgar F. Codd Innovations Award

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

Dr. Umeshwar Dayal is the recipient of the 2010 SIGMOD Edgar F. Codd Innovation Award for a succession of pioneering, influential contributions in distributed heterogeneous databases, high-performance active databases, generalized transitive closure, transaction models for long-running activities, and business process discovery, among other topics.

Details: Umeshwar Dayal has a track record of 30 years of research accomplishments in data management and has made a succession of fundamental contributions to the field. His research in the mid-1980's on Multibase, the world's first large-scale heterogeneous distributed database system, paved the way for research over the past two decades at universities and industrial research labs (HP Labs, IBM Research, and elsewhere) in data integration and federated databases. PROBE, one of the first extensible database management systems, made significant contributions to the field in several areas, including object and semantic data modeling, spatial, temporal and recursive query processing, and system architecture. HiPAC was one of the leading active database systems of its time, and the only one to focus on the needs of real-time applications, resulting in an innovative transaction model. The event-condition-action (ECA) rule model introduced in HiPAC has now been widely adopted in reactive computing systems, complex-event-processing systems, and distributed middleware.

Dayal also has significant results in query-processing research, with particularly strong contributions to the processing of multi-database queries, spatial queries, and recursive queries. The approach described in his VLDB 1987 paper on unnesting SQL queries was later adopted in at least five commercial products. In addition, Dayal performed pioneering work in long-duration transactions, business-process management, and database design. In particular, he pioneered the field of business-process intelligence, which combines data warehousing, data mining, analytics and optimization techniques to monitor, control, analyze, and optimize business processes.

Over 160 research papers and over 25 patents testify to Dayal's innovation and productivity. In 2001, he received the prestigious 10-year best paper award from VLDB for his paper on a transactional model for long-running activities. He is an HP Fellow, recognized for career contributions that "caused substantial change" in the state of the art while also improving HP products. In addition to his many innovative technical contributions, he has a distinguished record of service to our community, as a member of the editorial board of major journals (including ACM TODS and VLDB), chairing conferences (including SIGMOD, VLDB and ICDE), serving on boards and steering committees (VLDB, IEEE TC on Electronic Commerce, SIAM International Conference on Data Mining among them), and mentoring junior colleagues and young researchers.

SIGMOD Contributions Award

for outstanding and sustained services to and promotion of the database field through activities such as education, conference organization, journals, standards, and research funding.

Dr. David Lomet is the recipient of the 2010 SIGMOD Contribution Award for his outstanding leadership as the Editor-In-Chief of the IEEE Data Engineering Bulletin, a key forum for dissemination of emerging ideas in academia and industry. Lomet has been a key figure in our field, holding many additional leadership roles and demonstrating in each his dedication to service and to our community.

Details: By awarding David Lomet the ACM Contributions Award, we recognize his outstanding contributions to our community in leading the IEEE Data Engineering Bulletin for nearly 20 years and thereby creating a collection of timely articles of great value. He has almost single-handedly driven the IEEE Data Engineering Bulletin, providing our community with a constant stream of special issues, assembled by world-class invited guest editors. This service has been a wonderful benefit to the field, as each issue has provided a "root node" into key projects, both academic and industrial, and into the research literature related to the topic of the issue. These articles and issues have thus provided a "fast path" to see what's happening in an area as well as a way to make sure industrial highlights are brought to the attention of academics and vice versa. Not only has Lomet run the Bulletin, but he also negotiated with IEEE Computer Society and authors to make issues of the Bulletin available on CDROM via the SIGMOD DiSC, and later, to digitize the entire set of issues from 1977 on so that they can be available online to all.

Lomet has made significant contributions to our field through service in other roles as well. He was a Co-PC Chair of VLDB and he is currently on the VLDB Board of Trustees. He has been both a Co-PC Chair and a Co-General Chair of IEEE Data Engineering Conferences, and served as a member of the Steering Committee of the IEEE Technical Committee on Data Engineering from 2004-2009. He sets a high standard of service for our community.

SIGMOD Test-of-Time Award

for the paper from the 1999 SIGMOD Conference that has had the most impact (research, products, methodology) over the intervening decade.

NiagaraCQ: A Scalable Continuous Query System for Internet Databases
Jianjun Chen, David J. DeWitt (University of Wisconsin, now Microsoft), Feng Tian (University of Wisconsin, now VMWare), Yuan Wang (University of Wisconsin, now Microsoft)

This paper from the SIGMOD 2000 Conference bridged from the world of continuous, or standing, queries against a changing stored database, to stream processing systems. NiagaraCQ was a pioneering system, the first to address the problem of the millions of overlapping queries that would need to be supported in a truly internet-scale system. It used relational-style operators to optimize a given set of continuous queries. Similar frameworks appeared in subsequent studies of stream databases, sensor databases, information delivery systems, and complex-event-processing (CEP) systems. The idea of dynamic optimization of continuous queries leveraging database operators (including dynamic query grouping and split) became a baseline for modern streaming data platforms. In summary, this paper helped open the new field of high-performance systems for continuous query processing, and was a strong force in shaping the following generations of stream processing systems.

Original abstract: *Continuous queries are persistent queries that allow users to receive new results when they become available. While continuous query systems can transform a passive web into an active environment, they need to be able to support millions of queries due to the scale of the Internet. No existing systems have achieved this level of scalability. NiagaraCQ addresses this problem by grouping continuous queries based on the observation that many web queries share similar structures. Grouped queries can share the common computation, tend to fit in memory and can reduce the I/O cost significantly. Furthermore, grouping on selection predicates can eliminate a large number of unnecessary query invocations. Our grouping technique is distinguished from previous group optimization approaches in the following ways. First, we use an incremental group optimization strategy with dynamic re-grouping. New queries are added to existing query groups, without having to regroup already installed queries. Second, we use a query-split scheme that requires minimal changes to a general-purpose query engine. Third, NiagaraCQ groups both change-based and timer-based queries in a uniform way. To insure that NiagaraCQ is scalable, we have also employed other techniques including incremental evaluation of continuous queries, use of both pull and push models for detecting heterogeneous data source changes, and memory caching. This paper presents the design of NiagaraCQ system and gives some experimental results on the system's performance and scalability.*

SIGMOD Best Paper Award

FAST: Fast Architecture Sensitive Tree Search on Modern CPUs and GPUs

Changkyu Kim, Jatin Chhugani, Nadathur Satish (Intel), Eric Sedlar (Oracle), Anthony Nguyen (Intel), Tim Kaldewey (Oracle), Victor Lee (Intel), Scott Brandt (University of California, Santa Cruz), Pradeep Dubey (Intel)

Citation by the Awards Committee: This paper presents FAST, a layout for an in-memory binary tree index that is well-suited for state-of-the-art CPU and GPU architectures. The layout and associated search methods take advantage of SIMD instructions and thread-level parallelism (TLP). FAST also accounts for cache-line sizes and hides cache-miss and TLB-miss latency by processing many outstanding queries simultaneously (with software pipelining and TLP). The paper shows that with all these optimizations, search on GPU is compute bound and search on a CPU is bandwidth bound. To optimize the latter further, the paper presents a key-compression scheme, which also takes advantage of SIMD instructions, to alleviate bandwidth limits and handle larger keys. Experiments show how CPU and GPU perform on trees with different sizes, how many concurrent queries are needed to achieve their peak throughput, and how compression can improve search performance. This paper is an excellent research contribution that provides an end-to-end system design and associated algorithms and techniques to develop a complete solution that leverages the underlying hardware architecture. Given the modular structure of the overall design, the solution can easily be adapted to future architectures.

PODS Best Paper Award

An Optimal Algorithm for the Distinct Elements Problem

Daniel M. Kane, Jelani Nelson and David P. Woodruff

Abstract: We give the first optimal algorithm for estimating the number of distinct elements in a data stream, closing a long line of theoretical research on this problem begun by Flajolet and Martin in their seminal paper in FOCS 1983. This problem has applications to query optimization, Internet routing, network topology, and data mining. For a stream of indices in {1,...,n}, our algorithm computes a (1 +/- eps)-approximation using an optimal O(eps^{-2} + log(n)) bits of space with 2/3 success probability, where 0<eps<1 is given. This probability can be amplified by independent repetition. Furthermore, our algorithm processes each stream update in O(1) worst-case time, and can report an estimate at any point midstream in O(1) worst-case time, thus settling both the space and time complexities simultaneously. We also give an algorithm to estimate the Hamming norm of a stream, a generalization of the number of distinct elements, which is useful in data cleaning, packet tracing, and database auditing. Our algorithm uses nearly optimal space, and has optimal O(1) update and reporting times.

The ACM PODS Alberto O. Mendelzon Test-of-Time Award 2010

In 2007, the PODS Executive Committee decided to establish a Test-of-Time Award, named after the late Alberto O. Mendelzon, in recognition of his scientific legacy, and his service and dedication to the database community. Mendelzon was an international leader in database theory, whose pioneering and fundamental work has inspired and influenced both database theoreticians and practitioners, and continues to be applied in a variety of advanced settings. He served the database community in many ways; in particular, he served as the General Chair of the PODS conference, and was instrumental in bringing together the PODS and SIGMOD conferences. He also was an outstanding educator, who guided the research of numerous doctoral students and postdoctoral fellows. The Award is to be awarded each year to a paper or a small number of papers published in the PODS proceedings ten years prior, that had the most impact (in terms of research, methodology, or transfer to practice) over the intervening decade. The decision was approved by SIGMOD and the ACM. The funds for the Award were contributed by IBM Toronto. After careful consideration, the Award Committee for 2010 has decided to select the following papers as the award winners for 2010:

Typechecking for XML Transformers

Tova Milo, Dan Suciu, and Victor Vianu

Integrity Constraints for XML

Wenfei Fan and Jerome Simeon

SIGMOD TUTORIALS

SIGMOD Tutorial 1: Mining Knowledge from Databases: An Information Network Analysis Approach

Tuesday, 10:30 – 12:00 and 13:30 – 15:00
Location: Theory

Jiawei Han (University of Illinois, Urbana-Champaign), Yizhou Sun (University of Illinois, Urbana-Champaign), Xifeng Yan (UC Santa Barbara), and Philip S. Yu (University of Illinois, Chicago)

Most people consider a database is merely a data repository that supports data storage and retrieval. Actually, a database contains rich, inter-related, multi-typed data and information, forming one or a set of gigantic, interconnected, heterogeneous information networks. Much knowledge can be derived from such information networks if we systematically develop an effective and scalable database-oriented information network analysis technology. In this tutorial, we introduce database-oriented information network analysis methods and demonstrate how information networks can be used to improve data quality and consistency, facilitate data integration, and generate interesting knowledge. This tutorial presents an organized picture on how to turn a database into one or a set of organized heterogeneous information networks, how information networks can be used for data cleaning, data consolidation, and data qualify improvement, how to discover various kinds of knowledge from information networks, how to perform OLAP in information networks, and how to transform database data into knowledge by information network analysis. Moreover, we present interesting case studies on real datasets, including DBLP and Flickr, and show how interesting and organized knowledge can be generated from database-oriented information networks.

Jiawei Han is a Professor in the Department of Computer Science at the University of Illinois. He has served on many program committees of the major international conferences in the fields of data mining and database systems, and also served or is serving on the editorial boards for Data Mining and Knowledge Discovery, IEEE Transactions on Knowledge and Data Engineering, Journal of Computer Science and Technology, and Journal of Intelligent Information Systems. He is the founding Editor-in-Chief of ACM Transactions on Knowledge Discovery from Data (TKDD). Jiawei has received IBM Faculty Awards, HP Innovation Award, the Outstanding Contribution Award at the International Conference on Data Mining (2002), ACM SIGKDD Innovation Award (2004), IEEE Computer Society Technical Achievement Award (2005), and IEEE CS W. Wallace McDowell Award (2009). He is a Fellow of ACM and IEEE. He is currently the Director of Information Network Academic Research Center (INARC) supported by the Network Science-Collaborative Technology Alliance (NS-CTA) program of U.S. Army Research Lab. His book "Data Mining: Concepts and Techniques" (Morgan Kaufmann) has been used worldwide as a textbook.

Yizhou Sun is Ph.D. candidate at the University of Illinois at Urbana-Champaign. She received her M.S. on Signal and Information Processing in 2007 and B.S. degrees on Computer Science and Statistics in 2004, all from Peking Univ. She has developed RankClus and NetClus algorithms and is one of the team members developing a system, iNextCube, for information network analysis on multidimensional text databases.

Xifeng Yan is an assistant professor at the University of California at Santa Barbara, and holds the Venkatesh Narayanamurti Chair in Computer Science. He received a PhD degree in Computer Science from the University of Illinois at Urbana-Champaign in 2006. He was a research staff member at the IBM T. J. Watson Research Center between 2006 and 2008. Dr. Yan's research interests include data mining, databases, and bioinformatics. He has filed 6 patents and published more than 50 papers in refereed journals and conferences. His works are extensively referenced, with over 3,000 citations per Google Scholar. Dr. Yan received the 2007 ACM SIGMOD Doctoral Dissertation Runner-Up Award for his work in graph mining and graph data management.

Philip S. Yu is a Professor in the Department of Computer Science at the University of Illinois at Chicago and also holds the Wexler Chair in Information and Technology. He spent most of his career at IBM Thomas J. Watson Research Center and was manager of the Software Tools and Techniques group. Dr. Yu is a Fellow of the ACM and the IEEE. He is an associate editor of ACM Transactions on Knowledge Discovery from Data and serves on the steering committee of IEEE International Conference on Data Mining. Previously, he served as the Editor-in-Chief of IEEE Transactions on Knowledge and Data Engineering (2001-2004) and an associate editor for ACM Transactions of the Internet Technology (2000-2010). He was a member of the IEEE Data Engineering steering committee. He received an IEEE Region 1 Award for promoting and perpetuating numerous new electrical engineering concepts. His research interests include data mining, privacy, data stream, database systems and web technologies. Dr. Yu has published more than 570 papers in refereed journals and conferences. He holds or has applied for more than 300 US patents. Dr. Yu was an IBM Master Inventor. He received his PhD degree from Stanford University.

SIGMOD Tutorial 2: Database Systems Research on Data Mining
Tuesday, 15:30 – 17:00
Location: Theory

Carlos Ordonez (University of Houston) and Javier Garcia-Garcia (UNAM University, Mexico)

Data mining remains a broad and challenging problem in database systems. We present a review of processing alternatives, storage mechanisms, algorithms, data structures and optimizations that enable data mining on large data sets. We focus on the computation of several well-known multidimensional statistical and machine learning models. We pay particular attention to SQL, together with User-Defined Functions, and MapReduce as two competing and complementary technologies for large-scale processing. We conclude with a summary of solved major problems and open research issues.

Carlos Ordonez got a Ph.D. degree in Computer Science from the Georgia Institute of Technology, USA, in 2000. Dr Ordonez worked six years extending the Teradata DBMS with advanced data mining techniques to analyze large databases. Dr. Ordonez is currently an Assistant Professor at the University of Houston. His research is centered on the integration of machine learning and statistical techniques into database systems to analyze large data sets as well as their application to scientific problems. Dr Ordonez's research has produced over 60 papers, over 500 citations and has been funded by NSF.

Javier Garcia-Garcia received a PhD degree in computer science from UNAM University, Mexico. Dr Garcia-Garcia worked managing IT departments at several public and private institutions in Mexico for more than 20 years. Currently, he is a professor at the UNAM University, where he conducts research on relational database systems, focusing on data quality and data mining. He is a member of ACM and IEEE.

SIGMOD Tutorial 3: Information Theory for Data Management
Wednesday, 8:30 – 10:00 and 10:30 – 12:00
Location: Theory

Suresh Venkatasubramanian (AT&T), and Divesh Srivastava

We explore the use of information theory as a tool to express and quantify notions of information content and information transfer for representing and analyzing data, using examples from database design, data integration and data anonymization. We also examine the computational challenges associated with information-theoretic primitives, indicating how they might be computed efficiently.

Divesh Srivastava is the head of the Database Research Department at AT&T Labs-Research. He received his Ph.D. from the University of Wisconsin, Madison, and his B.Tech from the Indian Institute of Technology, Bombay. His research interests lie in data quality, data streams and data privacy.

Suresh Venkatasubramanian is the John and Marva Warnock Assistant Professor at the University of Utah. He received his Ph.D. from Stanford University, and his B.Tech from the Indian Institute of Technology, Kanpur. His research interests lie in algorithms, computational geometry, and large data mining and analysis.

SIGMOD Tutorial 4: Enterprise Information Extraction: Recent Developments and Open Challenges
Thursday, 10:30 – 12:00 and 1:30 – 3:00
Location: Theory

Laura Chiticariu (IBM Research, Almaden), Yunyao Li (IBM Research, Almaden), Sriram Raghavan (IBM Research, Almaden), and Frederick Reiss (IBM Research, Almaden)

Information extraction (IE)—the problem of extracting structured information from unstructured text — has become an increasingly important topic in recent years. A SIGMOD 2006 tutorial outlined challenges and opportunities for the database community to advance the state of the art in

information extraction, and posed the following grand challenge: “Can we build a System R for information extraction?”

Our tutorial gives an overview of progress the database community has made towards meeting this challenge. In particular, we start by discussing design requirements in building an enterprise IE system. We then survey recent technological advances towards addressing these requirements, broadly categorized as: (1) Languages for specifying extraction programs in a declarative way, thus allowing database-style performance optimizations; (2) Infrastructure needed to ensure scalability, and (3) Development support for enterprise IE systems. Finally, we outline several open challenges and opportunities for the database community to further advance the state of the art in enterprise IE systems. The tutorial is intended for students and researchers interested in information extraction

Laura Chiticariu is a Research Staff Member at IBM Research – Almaden. She received her Ph.D from U.C. Santa Cruz in 2008. Her current research focuses on using provenance to improve developmental support in information extraction systems.

Yunyao Li is a Research Staff Member at IBM Research – Almaden. She received her Ph.D from the University of Michigan, Ann Arbor in 2007. Her research focuses on improving information accessibility for enterprise applications.

Sriram Raghavan manages the Search and Analytics Group at IBM Research – Almaden. He has been involved in developing and building enterprise search and information extraction applications for the past 6 years, since graduating from Stanford University in 2004.

Frederick Reiss is a Research Staff Member at IBM Research – Almaden. He received his Ph.D. from U.C. Berkeley in 2006. His research focuses on improving the scalability of text analytics in enterprise applications.

PODS PAPER ABSTRACTS

PODS Session 1: PODS Opening and Keynote

Chair: Jan Paredaens (University of Antwerp)
Monday 8:30-9:45

Datalog Redux: Experience and Conjecture
Joe Hellerstein (UC Berkeley)

Datalog was a foundational topic in the early years of PODS, despite skepticism from practitioners about its relevance. This has been changing in recent years, with unlikely champions exploring and promoting its use as a practical basis for programming in a wide variety of application domains. We reflect on our use of Datalog to build systems of significant complexity for both networking and cloud computing infrastructure. Based on that experience, we present conjectures regarding next-generation programming languages, and the role that database theory could play in their development.

PODS Session 2: Query Languages

Chair: Tova Milo (Tel Aviv University)
Monday 10:15-11:45

Expressive Languages for Path Queries over Graph-Structured Data
Pablo Barceló (University of Chile), Carlos Hurtado (Adolfo Ibanez University), Leonid Libkin (University of Edinburgh), and Peter Wood (University of London)

For many problems arising in the setting of graph querying (such as finding semantic associations in RDF graphs, exact and approximate pattern matching, sequence alignment, etc.), the power of standard languages such as the widely studied conjunctive regular path queries (CRPQs) is insufficient in at least two ways. First, they cannot output paths and second, more crucially, they cannot express relations among paths. We thus propose a class of extended CRPQs, called ECRPQs, which add regular relations on tuples of paths, and allow path variables in the heads of queries. We provide several examples of their usefulness in querying graph structured data, and study their properties. We analyze query evaluation and representation of tuples of paths in the output by means of automata. We present a detailed analysis of data and combined complexity of queries, and look at restrictions that lower the complexity of ECRPQs to that of relational conjunctive queries. We study the containment problem, and look at further extensions with first-order features, and with non-regular relations that express arithmetic properties of paths, based on the lengths and numbers of occurrences of labels.

Transducing Markov Sequences
Benny Kimelfeld (IBM Almaden) and Christopher Re (University of Wisconsin)

A Markov sequence is a basic statistical model representing uncertain sequential data, and it is used within a plethora of applications, including speech recognition, image processing, computational biology, radio-frequency identification (RFID), and information extraction. The problem of querying a Markov sequence is studied under the conventional semantics of querying a probabilistic database, where queries are formulated as finite-state transducers. Specifically, the complexity of two main problems is analyzed. The first problem is that of computing the confidence (probability) of an answer. The second is the enumeration of the answers in the order of decreasing confidence (with the generation of the top-k answers as a special case), or in an approximate order thereof. In particular, it is shown that enumeration in any sub-exponential-approximate order is generally intractable (even for some fixed transducers), and a matching upper bound is obtained through a proposed heuristic. Due to this hardness, a special consideration is given to restricted (yet common) classes of transducers that extract matches of a regular expression (subject to prefix and suffix constraints), and it is shown that these classes are, indeed, significantly more tractable.

Positive Higher-Order Queries
Michael Benedikt (Oxford University), Gabriele Puppis (Oxford University), and Huy Vu (Oxford University)

We investigate a higher-order query language that embeds positive relational operators within the simply-typed λ-calculus. Our language allows one to succinctly define ordinary positive relational algebra queries (conjunctive queries and unions of conjunctive queries), and also supports the definition of second-order query functionals, allowing the transformation of CQs and UCQs in a generic (i.e. syntax-independent) way. We investigate the equivalence and containment problems for this calculus, which subsumes traditional CQ/UCQ containment. Equivalence of query functionals is defined via the notion of ``higher-order containment'', which states that the input to each functional is contained in its output. These notion of containment and equivalence depend on the class of (ordinary relational algebra) queries considered. We show that containment and equivalence are decidable when query variables are restricted to range over positive relational algebra, and identify the precise complexity of the problem. We identify classes of functionals where containment is tractable. We also provide bounds on the containment problem when functionals act over other classes of queries.

PODS Session 3: Awards

Chair: Jianwen Su (Univ. of California, Santa Barbara)
Monday 13:15-14:30

The ACM PODS Alberto O. Mendelzon Test-of-Time Award 2010:
Typechecking for XML Transformers
Tova Milo (Tel Aviv University), Dan Suciu (University of Washington), and Victor Vianu (UC San Diego)

The ACM PODS Alberto O. Mendelzon Test-of-Time Award 2010:
Integrity Constraints for XML
Wenfei Fan (University of Edinburgh) and Jerome Simeon (IBM T.J. Watson)

The first paper studied the problem of checking whether or not an XML transformation (e.g., specified in XSLT and in other languages) is well typed: for every input document of a given DTD, the result document always conforms to another specified DTD. The problem is essential for manipulating XML documents. The main result of the paper is that the problem is decidable. A key ingredient in obtaining this result was the introduction and study of a new tree-transducer model with pebbles that serves as an abstraction of XML transformations.

The second paper investigated integrity constraints for XML DTDs, including keys, foreign keys, inverse constraints, and inclusion dependencies. The technical results concern the implication and finite implication problems for the constraints. Clearly, integrity constraints and, in particular, the types studied in this paper, are an essential part of XML data modeling in a variety of contexts, including data management and software engineering. The undecidability, decidability, and complexity results provide a guidance and basis for dealing with integrity constraints in practice.

Both papers are extensively cited in the literature, and both had a major influence on the methodology and direction of subsequent research on XML data modeling and management. Hence, the committee has found both to be worthy of this Award.

Best Paper Award:
An Optimal Algorithm for the Distinct Elements Problem
Daniel M. Kane (Harvard University), Jelani Nelson (MIT), and David P. Woodruff (IBM Almaden)

We give the first optimal algorithm for estimating the number of distinct elements in a data stream, closing a long line of theoretical research on this problem begun by Flajolet and Martin in their seminal paper in FOCS 1983. This problem has applications to query optimization, Internet routing,

network topology, and data mining. For a stream of indices in {1,...,n}, our algorithm computes a (1 +/- eps)-approximation using an optimal O(eps^{-2} + log(n)) bits of space with 2/3 success probability, where 0<eps<1 is given. This probability can be amplified by independent repetition. Furthermore, our algorithm processes each stream update in O(1) worst-case time, and can report an estimate at any point midstream in O(1) worst-case time, thus settling both the space and time complexities simultaneously. We also give an algorithm to estimate the Hamming norm of a stream, a generalization of the number of distinct elements, which is useful in data cleaning, packet tracing, and database auditing. Our algorithm uses nearly optimal space, and has optimal O(1) update and reporting times.

Regular Paper: Understanding Cardinality Estimation using Entropy Maximization
Christopher Re (University of Wisconsin) and Dan Suciu (University of Washington)

Cardinality estimation is the problem of estimating the number of tuples returned by a query; it is a fundamentally important task in data management today, used in query optimization, progress estimation, and resource provisioning. We study cardinality estimation in a fundamental and principled framework: given a set of statistical assertions about the number of tuples returned by a fixed set of queries, predict the number of tuples returned by a new query. We model this problem using the probability space, over the possible worlds of the databases, that satisfies all statistical assertions and maximizes the entropy. We call this the Entropy Maximization model for statistics (MaxEnt). In this paper we develop the mathematical techniques needed to use the MaxEnt model for predicting the cardinality of conjunctive queries.

PODS Session 4: Tutorial I

Chair: Marcelo Arenas (Pontifical Catholic Univ. of Chile)
Monday 14:45-15:45

From Information to Knowledge: Harvesting Entities and Relationships from Web Sources
Gerhard Weikum (Max Planck Institute for Informatics) and Martin Theobald (Max Planck Institute for Informatics)

There are major trends to advance the functionality of search engines to a more expressive semantic level. This is enabled by the advent of knowledge-sharing communities such as Wikipedia and the progress in automatically extracting entities and relationships from semistructured as well as natural-language Web sources. Recent endeavors of this kind include DBpedia, EntityCube, KnowItAll, ReadTheWeb, and our own YAGO-NAGA project (and others). The goal is to automatically construct and maintain a comprehensive knowledge base of facts about named entities, their semantic classes, and their mutual relations as well as temporal contexts, with high precision and high recall. This tutorial discusses state-of-the-art methods, research opportunities, and open challenges along this avenue of knowledge harvesting. It also addresses issues of querying the knowledge base and ranking answers.

PODS Session 5: Streams and Query Processing

Chair: Mikhail Atallah (Purdue University)
Monday 16:15-18:15

Optimal Sampling From Distributed Streams
Graham Cormode (AT&T Labs Research), Ke Yi (HKUST), Qin Zhang (HKUST), and S. Muthukrishnan (Rutgers University)

A fundamental problem in data management is to draw a sample of a large data set, for approximate query answering, selectivity estimation, and query planning. With large, streaming data sets, this problem becomes particularly difficult when the data is shared across multiple distributed sites. The challenge is to ensure that a sample is drawn uniformly across the union of the data while minimizing the communication needed to run the protocol and track parameters of the evolving data. At the same time, it is also necessary to make the protocol lightweight, by keeping the space and time costs low for each participant. In this paper, we present communication-efficient protocols for

sampling (both with and without replacement) from k distributed streams. These apply to the case when we want a sample from the full streams, and to the sliding window cases of only the W most recent items, or arrivals within the last w time units. We show that our protocols are optimal, not just in terms of the communication used, but also that they use minimal or near minimal (up to logarithmic factors) time to process each new item, and space to operate.

Incremental Query Evaluation Revisited
Christoph Koch (Cornell University)

This paper approaches the incremental view maintenance problem from an algebraic perspective. We start by defining a ring of databases, which, by supporting distributivity and an additive inverse, admits the representation of expressions by polynomials and is closed under taking deltas. We embed this ring into a query calculus that connects terms and formulas via aggregation and allows to express powerful SQL aggregate queries with nesting. Queries of this calculus are infinitely differentiable, and the k-th delta of a polynomial of degree k without nesting is purely a function of the update, not of the database. The main technical result is that, for non-nested queries, each individual aggregate value can be incrementally maintained using a constant amount of work. This is not possible for nonincremental evaluation.

Fast Manhattan Sketches in Data Streams
Jelani Nelson (MIT) and David P. Woodruff (IBM Almaden)

The L1-distance, also known as the Manhattan or taxicab distance, between two vectors x, y in R^n is sum_{i=1}^n |x_i-y_i|. Approximating this distance is a fundamental primitive on massive databases, with applications to clustering, nearest neighbor search, network monitoring, regression, sampling, and support vector machines. We give the first 1-pass streaming algorithm for this problem in the turnstile model with O~(eps^{-2}) space and O~(1) update time, where the O~(.) notation hides polylogarithmic factors. All previous algorithms either required Omega~(eps^{-3}) space or Omega~(eps^{-2}) update time and/or could not work in the turnstile model (i.e., support an arbitrary number of updates to each coordinate). Our bounds are optimal up to O~(1) factors.

Semantic Query Optimization in the Presence of Types
Michael Meier (University of Freiburg), Michael Schmidt (University of Freiburg), Fang Wei (University of Freiburg), and Georg Lausen (University of Freiburg)

Both semantic and type-based query optimization rely on the idea that queries often exhibit non-trivial rewritings if the state space of the database is restricted. Despite their close connection, these two problems to date have always been studied separately. We present a unifying, logic-based framework for query optimization in the presence of data dependencies and type information. It builds upon the classical chase algorithm and extends existing query minimization techniques to considerably larger classes of queries and dependencies. In particular, our setting requires chasing conjunctive queries (possibly with union and negation) in the presence of dependencies containing negation and disjunction. We study the applicability of the chase in this setting, develop novel conditions that guarantee its termination, identify fragments for which minimal query computation is always possible (w.r.t. a generic cost function), and investigate the complexity of related decision problems.

PODS Session 6: Privacy
Chair: Graham Cormode (AT&T Labs Research)
Tuesday 10:30-12:00

Optimizing Linear Counting Queries under Differential Privacy
Chao Li (University of Massachusetts Amherst), Michael Hay (University of Massachusetts Amherst), Vibhor Rastogi (University of Washington), Gerome Miklau (University of Massachusetts Amherst), and Andrew McGregor (University of Massachusetts Amherst)

Differential privacy is a robust privacy standard that has been successfully applied to a range of data

analysis tasks. Despite much recent work, optimal strategies for answering a collection of correlated queries are not known. We study the problem of devising a set of strategy queries, to be submitted and answered privately, that will support the answers to a given workload of queries. We propose a general framework in which query strategies are formed from linear combinations of counting queries, and we describe an optimal method for deriving new query answers from the answers to the strategy queries. Using this framework we characterize the error of strategies geometrically, and we propose solutions to the problem of finding optimal strategies.

Universally Optimal Privacy Mechanisms for Minimax Agents
Mangesh Gupte (Rutgers University) and Mukund Sundararajan (Google)

A scheme that publishes aggregate information about sensitive data must resolve the trade-off between utility to information consumers and privacy of the database participants. Differential privacy~\cite{BigBang} is a well-established definition of privacy---this is a universal guarantee against all attackers, whatever their side-information or intent. In this paper, we present a comparably universal treatment of utility based on the standard minimax rule from decision theory (in contrast to the utility model in~\cite{Bayesian}, which is Bayesian). In our model, information consumers are minimax decision makers, each possessing some side-information about the query, and each endowed with a loss-function which models their tolerance to inaccuracies. Further, information consumers are rational in the sense that they actively combine information from the mechanism with their side-information in a way that minimizes their loss. Under this assumption of rational behavior, we show that for every fixed count query, a certain geometric mechanism is universally optimal for all minimax information consumers. Additionally, our solution makes it possible to release query results at multiple levels of privacy in a collusion resistant manner.

Towards an Axiomatization of Statistical Privacy and Utility
Daniel Kifer (Penn State University) and Bing-Rong Lin (Penn State University)

``Privacy" and ``utility" are words that frequently appear in the literature on statistical privacy. But what do those words really mean? In recent years, many problems with intuitive notions of privacy and utility have been uncovered. Thus more formal notions of privacy and utility, which are amenable to mathematical analysis, are needed. In this paper we present our initial work on an axiomatization of privacy and utility. In particular, we study how these concepts are affected by randomized algorithms. Our analysis yields new insights into the construction of both privacy definitions and mechanisms that generate data according to such definitions. In particular, it characterizes a class of relaxations of differential privacy and shows that desirable outputs of a differentially private mechanism are best interpreted as certain graphs rather than query answers or synthetic data.

PODS Session 7: Tutorial 2 & Information Complexity & Missing Information

Chair: Jeffrey Naughton (Univ. of Wisconsin)
Tuesday 13:30-15:00

Tutorial 2: Information Complexity
T.S. Jayram (IBM Almaden)

The recent years have witnessed the overwhelming success of algorithms that operate on massive data. Several computing paradigms have been proposed for massive data set algorithms such as data streams, sketching, sampling etc. and understanding their limitations is a fundamental theoretical challenge. In this survey, we describe the information complexity paradigm that has proved successful in obtaining tight lower bounds for several well-known problems. Information complexity quantifies the amount of information about the inputs that must be necessarily propagated by any algorithm in solving a problem. We describe the key ideas of this paradigm, and highlight the beautiful interplay of techniques arising from diverse areas such as information theory, statistics and geometry.

Regular Paper: Capturing Missing Tuples and Missing Values

Floris Geerts (University of Edinburgh) and Wenfei Fan (University of Edinburgh)

Databases in real life are often neither entirely closed-world nor entirely open-world. Indeed, databases in an enterprise are typically partially closed, in which a part of the data is constrained by master data that contains complete information about the enterprise in certain categories. It has been shown that despite missing tuples, such a database may turn out to have complete information for answering a query. This paper studies partially closed databases from which both tuples and values may be missing. We specify such a database in terms of conditional tables constrained by master data, referred to as c-instances. We first propose three models to characterize whether a c-instance T is complete for a query Q relative to master data. That is, depending on how missing values in T are instantiated, the answer to Q in T remains unchanged when new tuples are added. We then investigate four problems, to determine (a) whether a given c-instance is complete for a query Q, (b) whether there exists a c-instance that is complete for Q relative to master data available, (c) whether a c-instance is a minimal-size database that is complete for Q, and (d) whether there exists a c-instance of a bounded size that is complete for Q. We establish matching lower and upper bounds on these problems for queries expressed in a variety of languages, in each of the three models for characterizing relative completeness.

PODS Sessions 8 & 9: Uncertainty in Databases

Chair: Benny Kimelfeld (IBM Almaden)

Tuesday 15:30-17:30

On the First-order Expressibility of Computing Certain Answers to Conjunctive Queries over Uncertain Databases

Jef Wijsen (University of Mons)

A natural way for capturing uncertainty in the relational model is by having relations that violate their primary key constraint, that is, relations in which distinct tuples agree on the primary key. A repair (or possible world) of a database is then obtained by selecting a maximal number of tuples without ever selecting two distinct tuples that have the same primary key value. For a Boolean query q, CERTAINTY(q) is the problem that takes as input a database db and asks whether q evaluates to true on every repair of db. We are interested in determining queries q for which CERTAINTY(q) is first-order expressible (and hence in the low complexity class AC0). For queries q in the class of conjunctive queries without self-join, we provide a necessary syntactic condition for first-order expressibility of CERTAINTY(q). For acyclic queries (in the sense of [Beeri et al. 83]), this necessary condition is also a sufficient condition. So we obtain a decision procedure for first-order expressibility of CERTAINTY(q) when q is acyclic and without self-join. We also show that if CERTAINTY(q) is first-order expressible, its first-order definition, commonly called (certain) first-order rewriting, can be constructed in a rather straightforward way.

Certain Answers for XML Queries

Claire David (University of Edinburgh), Leonid Libkin (University of Edinburgh), and Filip Murlak (University of Warsaw)

The notion of certain answers arises when one queries incompletely specified databases, e.g., in data integration and exchange scenarios, or databases with missing information. While in the relational case this notion is well understood, there is no natural analog of it for XML queries that return documents. We develop an approach to defining certain answers for such XML queries, and apply it in the settings of incomplete information and XML data exchange. We first revisit the relational case, and show how to present the key concepts related to certain answers in a new model-theoretic language. This new approach naturally extends to XML. We prove a number of generic, application-independent results about computability and complexity of certain answers produced by it. We then turn our attention to a pattern-based XML query language with trees as outputs, and present a technique for computing certain answers that relies on the notion of a basis of a set of trees. We show how to compute such bases for documents with nulls and for documents arising in data exchange scenarios, and provide complexity bounds. While in general complexity of query

answering in XML data exchange could be high, we exhibit a natural class of XML schema mappings for which not only query answering, but also many static analysis problems can be solved efficiently.

Computing Query Probability with Incidence Algebras

Nilesh Dalvi (Yahoo! Research), Karl Schnaitter (UC Santa Cruz), and Dan Suciu (University of Washington)

We describe an algorithm for evaluating queries over probabilistic databases using incidence algebras. The queries we consider are unions of conjunctive queries, and the probabilistic database are tuple-independent structures. Our algorithm runs in PTIME, on a subset of queries called "safe" queries. The algorithm is very simple, and easy to implement in practice, yet it is highly non-obvious. The role played by the incidence algebras is that it allows us to avoid computing subqueries that are provably hard.

On Probabilistic Fixpoint and Markov Chain Query Languages

Daniel Deutch (Tel Aviv University), Cristoph Koch (Cornell University), and Tova Milo (Tel Aviv University)

We study highly expressive query languages such as datalog, fixpoint, and while-languages on probabilistic databases. We generalize these languages such that computation steps (e.g. datalog rules) can fire probabilistically. We define two possible semantics for such query languages, namely inflationary semantics where the results of each computation step are added to the current database and non-inflationary queries that induce a random walk over database instances. We then study the complexity of exact and approximate query evaluation under these semantics.

PODS Session 10: Schema mappings and design

Chair: Jan Van den Bussche (Hasselt University)
Wednesday 8:30-10:00

Foundations of Schema Mapping Management

Marcelo Arenas (Pontifical Catholic University of Chile), Jorge Pérez (Pontifical Catholic University of Chile), Juan L. Reutter (University of Edinburgh), and Cristian Riveros (Oxford University)

A schema mapping is a specification that describes a relationship between data in two independent databases. In the last few years, a lot of attention has been paid to the specification and subsequent manipulation of schema mappings, a problem known as metadata management. There have been many achievements in this area, and semantics have been defined for operators such as composition and inverse. However, the absence of formal tools to compare schema mappings, in terms of their ability to transfer data and avoid storing redundant information, has hampered the development of foundations for more complex operators, as many of them involve these notions. In this paper, we address the problem of providing foundations for metadata management by developing a measure of the amount of information transferred by a schema mapping. More precisely, we first identify the necessary conditions that an order on the amount of information transferred by a schema mapping should satisfy, and then we propose an order <=_S that is provably the strictest relation satisfying these conditions. Second, we study some fundamental properties of <=_S and, in particular, we show its decidability for the class of mappings specified by CQ^{\neq}-TO-CQ dependencies, which includes the widely used class of source-to-target tuple-generating dependencies (st-tgds). Third, as a proof of concept, we show how the machinery developed for the order <=_S can be used to study some metadata management problems in the context of data exchange. In particular, we provide simpler proofs for some fundamental results regarding the inverse operator, and we give an effective characterization for the decidability of the well-known schema evolution problem. Fourth, we use all this machinery to study some more complex metadata management operators. In particular, we study the semantics of the extract operator, that intuitively captures the idea of upgrading a legacy schema, and we provide an algorithm for computing it for the class of mappings specified by st-tgds. Moreover, we also study the merge operator, that, as well as the extract operator, has been identified as fundamental for the development of a metadata management framework.

Schema Design for XML Repositories: Complexity and Tractability

Wim Martens (Technical University of Dortmund), Matthias Niewerth (Technical University of Dortmund), and Thomas Schwentick (Technical University of Dortmund)

Abiteboul et al. initiated the systematic study of distributed XML documents consisting of several logical parts, possibly located on different machines. The physical distribution of such documents immediately raises the following question: how can a global schema for the distributed document be broken up into local schemas for the different logical parts? The desired set of local schemas should guarantee that, if each logical part satisfies its local schema, then the distributed document satisfies the global schema. Abiteboul et al. proposed three levels of desirability for local schemas: local typing, maximal local typing, and perfect local typing. Immediate algorithmic questions are: (i) given a typing, determine whether it is local, maximal local, or perfect, and (ii) given a document and a schema, establish whether a (maximal) local or perfect typing exists. This paper improves the open complexity results in their work and initiates the study of (i) and (ii) for schema restrictions arising from the current standards: DTDs and XML Schemas with deterministic content models. The most striking result is that these restrictions yield tractable complexities for the perfect typing problem. Furthermore, an open problem in Formal Language Theory is settled: deciding language primality for deterministic finite automata is pspace-complete.

Simplifying XML Schema: Single-Type Approximations of Regular Tree Languages

Wouter Gelade (Hasselt University/Transnational University of Limburg), Tomasz Idziaszek (University of Warsaw), Wim Martens (Technical University of Dortmund), and Frank Neven (Hasselt University/Transnational University of Limburg)

XML Schema Definitions (XSDs) can be adequately abstracted by the single-type regular tree languages. It is well-known, that these form a strict subclass of the robust class of regular unranked tree languages. Sadly, in this respect, XSDs are not closed under the basic operations of union and set difference, complicating important task in schema integration and evolution. The purpose of this paper is to investigate how the union and difference of two XSDs can be approximated within the framework of single-type regular tree languages. We consider both lower and upper approximations. We also address the more general question of how to approximate an arbitrary regular tree language by an XSD.

PODS Session 11: Query Learning

Chair: Dan Suciu (University of Washington)
Wednesday 10:30-12:00

Characterizing Schema Mappings via Data Examples

Bogdan Alexe (UC Santa Cruz), Phokion Kolaitis (UC Santa Cruz and IBM Almaden), and Wang-Chiew Tan (UC Santa Cruz)

Schema mappings are high-level specifications that describe the relationship between two database schemas; they are considered to be the essential building blocks in data exchange and data integration, and have been the object of extensive research investigations. Since in real-life applications schema mappings can be quite complex, it is important to develop methods and tools for understanding, explaining, and refining schema mappings. A promising approach to this effect is to use "good" data examples that illustrate the schema mapping at hand. We develop a foundation for the systematic investigation of data examples and obtain a number of results on both their capabilities and limitations in explaining and understanding schema mappings. We focus on schema mappings specified by source-to-target tuple generating dependencies (s-t tgds) and investigate the following problem: which classes of s-t tgds can be "uniquely characterized" by a finite set of data examples? Our investigation begins by considering finite sets of positive and negative examples, which are arguably the most natural choice of data examples. However, we show that they are not powerful enough to yield interesting unique characterizations. We then consider finite sets of universal examples, where a universal example is a pair consisting of a source instance and a universal solution for that source instance. We unveil a tight connection between unique characterizations with universal examples and the existence of Armstrong bases (a relaxation of the

classical notion of Armstrong databases). On the positive side, we show that every schema mapping specified by LAV s-t tgds is uniquely characterized by a finite set of universal examples with respect to the class of LAV s-t tgds. Moreover, this positive result extends to the much broader classes of n-modular schema mappings, n a positive integer. Finally, we show that, on the negative side, there are schema mappings specified by GAV s-t tgds that are not uniquely characterized by any finite set of universal examples and negative examples with respect to the class of GAV s-t tgds (hence also with respect to the class of all s-t tgds).

Understanding Queries in a Search Database System

Ronald Fagin (IBM Almaden), Benny Kimelfeld (IBM Almaden), Yunyao Li (IBM Almaden), Sriram Raghavan (IBM Almaden), and Shivakumar Vaithyanathan (IBM Almaden)

It is well known that a search engine can significantly benefit from an auxiliary database, which can suggest interpretations of the search query by means of the involved concepts and their interrelationship. The difficulty is to translate abstract notions like concept and interpretation into a concrete search algorithm that operates over the auxiliary database. To surpass existing heuristics, there is a need for a formal basis, which is realized in this paper through the framework of a search database system, where an interpretation is identified as a parse. It is shown that the parses of a query can be generated in polynomial time in the combined size of the input and the output, even if parses are restricted to those having a nonempty evaluation. Identifying that one parse is more specific than another is important for ranking answers, and this framework captures the precise semantics of being more specific; moreover, performing this comparison between parses is tractable. Lastly, the paper studies the problem of finding the most specific parses. Unfortunately, this problem turns out to be intractable in the general case. However, under reasonable assumptions, the parses can be enumerated in an order of decreasing specificity, with polynomial delay and polynomial space.

A Learning Algorithm for Top-Down XML Transformations

Aurelien Lemay (University of Lille & INRIA), Sebastian Maneth (NICTA and UNSW), and Joachim Niehren (INRIA)

A classical result from learning theory says that for any regular language L there exists a characteristic sample set of positive and negative examples which allows to infer a representation of L. In fact, there are such sample sets of polynomial size in the size of the minimal DFA for L. This is proved through the Myhill-Nerode Theorem for regular languages. Here, a generalization of this result from string to trees and from languages to translations is given: for any deterministic top-down tree transformation there exists a sample set of polynomial size (with respect to the minimal transducer) which allows to infer the translation. Until now, only for string transducers and for simple relabeling tree transducers, similar results have been known. Learning of deterministic top-down tree transducers (DTOPs) is far more involved because: they can copy, delete, and permute their input subtrees. Thus, complex dependencies of labeled input to output paths need to be maintained by the algorithm. First, a Myhill-Nerode theorem is presented for DTOPs. The theorem is interesting on its own, and challenging to prove. This theorem is then used to construct a learning algorithm for DTOPs. Finally, we investigate how our result can be applied to XML transformations (e.g. XSLT programs). For this, a new (DTD based) encoding of unranked trees by ranked ones is presented. Over such encodings, DTOPs can realize a large class of practically interesting XML transformations, which cannot be realized over the widely used first-child next-sibling encoding.

PODS Session 12: Constraints & Indexing

Chair: Jan Hidders (Delft Univ. of Technology)

Wednesday 16:00-18:00

Cache-Oblivious Hashing

Rasmus Pagh (IT University of Copenhagen), Zhewei Wei (HKUST), Ke Yi (HKUST), and Qin Zhang (HKUST)

The hash table, especially its external memory version, is one of the most important index structures in large databases. Assuming a truly random hash function, it is known that in a standard external hash table, searching for a particular key only takes expected average $t_q=1+1/2^{\Omega(b)}$ disk accesses, where b is the disk block size. However, such near-perfect performance is achieved only when b is known and the hash table is particularly tuned for working with such a blocking. In this paper we study if it is possible to build a cache-oblivious hash table that works well with any blocking. Such a hash table will automatically perform well across all levels of the memory hierarchy and does not need any hardware specific tuning, an important feature in autonomous databases. We first show that linear probing, a classical collision resolution strategy for hash tables, can be easily made cache-oblivious but it only achieves $t_q = 1 + O(1/b)$. Then we demonstrate that it is possible to obtain $t_q = 1 + 1/2^{\Omega(b)}$, thus matching the cache-aware bound, if the following two conditions hold: (a) b is a power of 2; and (b) every block starts at a memory address divisible by b. Both conditions hold on a real machine, although they are not stated in the cache-oblivious model. Interestingly, we also show that neither condition is dispensable: if either of them is removed, the best obtainable bound is $t_q=1+O(1/b)$ with linear space, which is exactly what linear probing achieves.

Performance Guarantees for B-trees with Different-Sized Keys

Michael A. Bender (Stony Brook University and Tokutek, Inc.), Haodong Hu (Microsoft), and Bradley C. Kuszmaul (MIT and Tokutek, Inc.)

Most published descriptions of B-trees assume that all N keys have the same size. If keys all have size K and $f=B/K$ keys fit in a disk block, then the search cost is $O(\log_{f+1} N)$ block transfers. However, most production-quality B-trees support keys of different sizes. When keys have variable size, B-tree operations have no nontrivial performance guarantees. This paper bestows performance guarantees upon a B-tree when keys have different sizes. Technically, the paper focuses on atomic-key dictionaries. In an atomic-key dictionary (e.g., a B-tree), keys are stored and manipulated as atomic objects. Entire keys are stored in nodes or passed to the comparison function. Performance guarantees are parameterized by the average key size. The objective to give performance bounds in terms of the average key size matching those when all keys have the same size. The paper first gives static stuctures. If there are N keys with average size K, the search cost is $O(\lceil K/B \rceil \log_{1+\lceil B/K\rceil}N)$ expected memory transfers. The paper proves that it is not possible to transform these expected bounds into worst-case bounds. The cost to build this tree is $O(NK)$ operations and $O(NK/B)$ memory transfers if all keys are presented in sorted order. If not, the cost is the sorting cost. The paper next gives a dynamic structures. The amortized cost to insert a key $\key$ of arbitrary length and random rank is $O(\lceil K/B \rceil \log_{1+ \lceil B/K\rceil}N + \sizeof{\key}/B)$. More generally, the amortized cost to insert a key $\key$ of arbitrary length at an arbitrary rank is asymptotically the no more than the cost to search for $\key$. Thus, the search cost dominates the modification cost. Finally, the paper gives a dynamic-programming algorithm for constructing an optimal searching tree having the minimum expected cost. There already exist structures, e.g., the string B-tree, for storing strings, which are just different-size keys. String B-trees provably attain better asymptotic performance than atomic-key dictionaries. The string B-tree is not a solution for this paper, however, because it is not an atomic-key dictionary, and keys are not treated as atomic objects. In contrast, this paper's objective is to modify the traditional B-tree as little as possible, while giving performance guarantees when keys have different sizes.

When data dependencies over SQL tables meet the Logics of Paradox and S-3

Sven Hartmann (Clausthal University of Technology) and Sebastian Link (University of Wellington)

We study functional and multivalued dependencies over SQL tables with NOT NULL constraints. Under a no-information interpretation of nulls we develop tools for reasoning. We further show that in the absence of NOT NULL constraints the associated implication problem is equivalent to that in propositional fragments of Priest's paraconsistent Logic of Paradox. Subsequently, we extend the equivalences to Boolean dependencies and to the presence of NOT NULL constraints using Schaerf and Cadoli's S-3 logics where S corresponds to the union of attributes declared NOT NULL. The

findings also apply to Codd's interpretation ``value at present unknown'' utilizing a weak possible world semantics. Our results establish NOT NULL constraints as an effective mechanism to balance the expressiveness and tractability of consequence relations, and to control the degree by which the existing classical theory of data dependencies can be soundly approximated in practice.

The Power of Tree Projections: Local Consistency, Greedy Algorithms, and Larger Islands of Tractability

Gianluigi Greco (University of Calabria) and Francesco Scarcello (University of Calabria)

Enforcing local consistency is a well-known technique to simplify the evaluation of conjunctive queries. It consists of repeatedly taking the semijion between every pair of query atoms, until the procedure stabilizes. If some relation becomes empty, then the query has an empty answer. Otherwise, we cannot say anything in general, unless we have some information on the structure of the given query. In fact, a fundamental result in database theory states that the class of queries for which---on every database---local consistency entails global consistency is precisely the class of acyclic queries. In the last few years, several efforts have been spent to define structural decomposition methods isolating larger classes of nearly-acyclic queries, yet retaining the same nice properties as acyclic ones. In particular, it is known that queries having bounded (generalized) hypertree-width can be evaluated in polynomial time, and that this structural property is also sufficient to guarantee that local consistency solves the problem, as for acyclic queries. However, the precise power of such an approach was an open problem: Is it the case that bounded generalized hypertree-width is also a necessary condition to guarantee that local consistency entails global consistency? In this paper, we positively answer this question, and go beyond. Firstly, we precisely characterize the power of local consistency procedures in the more general framework of tree projections, where a query Q and a set V of views (i.e., resources that can be used to answer Q) are given, and where one looks for an acyclic hypergraph covering Q, and covered by V---all known structural decomposition methods are just special cases of this framework, defining their specific set of resources. We show that the existence of tree projections of certain subqueries is a necessary and sufficient condition to guarantee that local consistency allows the query to be answered efficiently, even without computing any tree projection. In particular, tight characterizations are given not only for the decision problem, but also when answers have to be computed. Secondly, we consider greedy tree-projections that are easy to compute, and we study how far they can be from arbitrary tree-projections, which are intractable in general. Finally, we investigate classes of instances not included in those having tree projections, and which can be easily recognized and define either new islands of tractability, or islands of quasi-tractability.

SIGMOD RESEARCH PAPER ABSTRACTS

SIGMOD Research Session 1: Advanced Query Processing

Chair: Walid Aref (Purdue University)

Tuesday, 10:30 – 12:00

Efficiently Evaluating Complex Boolean Expressions

Marcus Fontoura (Yahoo! Research), Suhas Sadanandan (Yahoo! Inc), Jayavel Shanmugasundaram, Sergei Vassilvitski, Erik Vee (Yahoo! Research), Srihari Venkatesan, Jason Zien (Yahoo! Inc)

The problem of efficiently evaluating a large collection of complex Boolean expressions –- beyond simple conjunctions and Disjunctive/Conjunctive Normal Forms (DNF/CNF) –- occurs in many emerging online advertising applications such as Advertising Exchanges and Automatic Targeting. The simple solution of normalizing complex Boolean expressions to DNF or CNF form, and then using existing methods for evaluating such expressions, is not always effective because of the exponential blow-up in the size of expressions due to normalization. We thus propose a novel method for evaluating complex expressions, which leverages existing techniques for evaluating leaf-level conjunctions, and then uses a bottom-up evaluation technique to only process the relevant parts of the complex expressions that contain the matching conjunctions. We develop two such bottom-up evaluation techniques, one based on Dewey IDs, and another based on mapping Boolean expressions to a one-dimensional space. Our experimental evaluation based on data obtained from an operational Advertising Exchange shows that the proposed techniques are efficient and scalable, both with respect to space usage as well as evaluation time.

How to ConQueR Why-Not Questions

Quoc Trung Tran, Chee-Yong Chan (National University of Singapore)

One useful feature that is missing from today's database systems is an explain capability that enables users to seek clarifications on unexpected query results. There are two types of unexpected query results that are of interest: the presence of unexpected tuples, and the absence of expected tuples (i.e., missing tuples). Clearly, it would be very helpful to users if they could pose follow-up why and why- not questions to seek clarifications on, respectively, unexpected and expected (but missing) tuples in query results. While the why questions can be addressed by applying established data provenance techniques, the problem of explaining the why-not questions has received very little attention. There are currently two explanation models proposed for why-not questions. The first model explains a missing tuple t in terms of modifications to the database such that t appears in the query result wrt the modified database. The second model explains by identifying the data manipulation operator in the query evaluation plan that is responsible for excluding t from the result. In this paper, we propose a new paradigm for explaining a why-not question that is based on automatically generating a refined query whose result includes both the original query's result as well as the user-specified missing tuple(s). In contrast to the existing explanation models, our approach goes beyond merely identifying the ``culprit'' query operator responsible for the missing tuple(s) and is useful for applications where it is not appropriate to modify the database to obtain missing tuples.

Call to Order: A Hierarchical Browsing Approach to Eliciting Users' Preference

Feng Zhao (National University of Singapore), Gautam Das (University of Texas at Arlington), Kian-Lee Tan, Anthony Tung (National University of Singapore)

Computing preference queries has received a lot of attentions in the database community. It is common that the user is unsure of his/her preference, so care must be taken to elicit the preference of the user correctly. In this paper, we propose to elicit the preferred ordering of a user by utilizing skyline objects as the representatives of the possible ordering. We introduce the notion of order-based representative skylines which selects representatives based on the orderings that they represent. To further facilitate preference exploration, a hierarchical clustering algorithm is applied

to compute a denogram on the skyline objects. By coupling the hierarchical clustering with visualization techniques, we allow users to refine their preference weight settings by browsing the hierarchy. Extensive experiments were conducted and the results validate the feasibility and the efficiency of our approach.

Boosting Spatial Pruning: On Optimal Pruning of MBRs

Tobias Emrich, Hans-Peter Kriegel, Peer Kroger, Matthias Renz, Andreas Zuefle (Ludwig-Maximilians-Universitaet Munchen)

Fast query processing of complex objects, e.g. spatial or uncertain objects, depends on efficient spatial pruning of the objects approximations, which are typically minimum bounding rectangles (MBRs). In this paper, we propose a novel effective and efficient criterion to determine the spatial topology between multi-dimensional rectangles. Given three rectangles R, A, and B in a multi-dimensional space, the task is to determine whether R is closer to A than to B. This \emph{domination} relation is used in many applications to perform spatial pruning. Traditional techniques apply spatial pruning based on minimal and maximal distance. These techniques however show significant deficiencies in terms of effectivity. We prove that our decision criterion correct, complete, and efficient to compute even for high dimensional databases. In addition, we tackle the problem of computing the number of objects dominating an object o. The challenge here is to incorporate objects that only partially dominate o. In this work we will show how to detect such partial domination topology by using a modified version of our decision criterion. We propose strategies for conservatively and progressively estimating the total number of objects dominating an object. Our experiments show that the new pruning criterion, albeit very general and widely applicable, significantly outperforms current state-of-the-art pruning criteria.

SIGMOD Research Session 2: Data Cleaning & Data Mining

Chair: Timos Sellis (National Technical University of Athens)
Tuesday, 10:30 – 12:00

Leveraging Spatio-Temporal Redundancy for RFID Data Cleansing

Haiquan Chen, Wei-shinn Ku (Auburn University), Haixun Wang (Microsoft Research, Asia), Min-Te Sun (National Central University, Taiwan)

Radio Frequency Identification (RFID) technologies are used in many applications for data collection. However, raw RFID readings are usually of low quality and may contain many anomalies. An ideal solution for RFID data cleansing should address the following issues. First, in many applications, duplicate readings (by multiple readers simultaneously or by a single reader over a period of time) of the same object are very common. The solution should take advantage of the resulting data redundancy for data cleaning. Second, prior knowledge about the readers and the environment (e.g., prior data distribution, false negative rates of readers) may help improve data quality and remove data anomalies, and a desired solution must be able to quantify the degree of uncertainty based on such knowledge. Third, the solution should take advantage of given constraints in target applications (e.g., the number of objects in a same location cannot exceed a given value) to elevate the accuracy of data cleansing. There are a number of existing RFID data cleansing techniques. However, none of them support all the aforementioned features. In this paper we propose a Bayesian inference based approach for cleaning RFID raw data. Our approach takes full advantage of data redundancy. To capture the likelihood, we design an n-state detection model and formally prove that the 3-state model can maximize the system performance. Moreover, in order to sample from the posterior, we devise a Metropolis-Hastings sampler with Constraints (MH-C), which incorporates constraint management to clean RFID raw data with high efficiency and accuracy. We validate our solution with a common RFID application and demonstrate the advantages of our approach through extensive simulations.

Sampling Dirty Data for Matching Attributes

Henning Koehler, Xiaofang Zhou , Shazia Sadiq (The University of Queensland), Yanfeng Shu, Kerry Taylor (CSIRO)

We investigate the problem of creating and analyzing samples of relational databases to find relationships between string-valued attributes. Our focus is on identifying attribute pairs whose value sets overlap, a pre-condition for allowing sensible joins. However, real-world data sets are often `dirty', especially when integrating data from different sources. To deal with this issue, we propose new similarity measures between sets of strings, which not only consider set similarity, but similarity between strings as well. To make the measures effective, we develop efficient algorithms for distributed sample creation and similarity computation. Test results show that for dirty data our measures (best used in combination) are more accurate for measuring value overlap than existing sample-based methods and scale well.

ERACER: A Database Approach for Statistical Inference and Data Cleaning

Chris Mayfield, Jennifer Neville, Sunil Prabhakar (Purdue University)

Real-world databases often contain both syntactic and semantic errors, in spite of integrity constraints and other safety measures incorporated into standard DBMSs. We present an iterative statistical framework, based on belief propagation and relational dependency networks, for inferring missing information and correcting such errors automatically. The key features of our approach are that (1) it uses an efficient approximate inference algorithm that is easily implemented in standard DBMSs and scales well to large databases sizes, and (2) it uses shrinkage techniques to accurately infer correct values even in the presence of corrupted or dirty data. Our framework performs these two tasks in an integrated manner, using standard SQL and user defined functions. We evaluate the method empirically on several synthetic and real-world data sets and compare to a baseline statistical method that uses Bayesian networks with exact inference. The results show that our algorithm achieves accuracy comparable to the baseline with respect to inferring missing values. Furthermore, our framework simultaneously identifies and corrects corrupted values with high precision.

Recsplorer: Recommendation Algorithms based on Precedence Mining

Aditya Parameswaran, Georgia Koutrika, Benjamin Bercovitz, Hector Garcia-Molina (Stanford University)

We study recommendations in applications where there are temporal patterns in the way items are consumed or watched. For example, a student who has taken the Advanced Algorithms course is more likely to be interested in Convex Optimization, but a student who has taken Convex Optimization need not be interested in Advanced Algorithms in the future. Similarly, a person who has purchased the Godfather I DVD on Amazon is more likely to purchase Godfather II sometime in the future (though it is not strictly necessary to watch/purchase Godfather I beforehand). We propose a precedence mining model that estimates the probability of future consumption based on past behavior. We then propose a suite of recommendation algorithms that exploit the precedence information. We evaluate our algorithms, as well as traditional recommendation ones, using a real course planning system. We use existing transcripts to evaluate how well the algorithms predict courses taken. In addition, we augment our experiments with a user study on the live system where users rate their recommendations.

SIGMOD Research Session 3: Graph Data & Querying

Chair: Graham Cormode (AT&T Research)

Tuesday, 10:30 – 12:00

TEDI: Efficient Shortest Path Query Answering on Graphs

Fang Wei (University of Freiburg)

Efficient shortest path query answering in large graphs is enjoying a growing number of applications, such as ranked keyword search in databases, social networks, ontology reasoning and bioinformatics. A shortest path query on a graph finds the shortest path from the given source and target vertices in the graph. Current techniques for efficient evaluation of such queries are based on the pre-computation of compressed Breadth First Search trees for the vertices. However, they suffer from drawbacks of scalability. To address these problems, we propose TEDI, an indexing and

query processing scheme for shortest path query answering. TEDI is based on the tree decomposition methodology. The graph is first decomposed into a tree in which each node (a.k.a. bag) contains more than one vertex from graph. The shortest paths are stored in such bags and these local paths together with the tree are the components of the index of the graph. Based on this index, a bottom-up operation can be executed to find the shortest path for any given source and target vertices. Our experimental results show that TEDI offers orders-of-magnitude performance improvement over existing approaches on the index construction time consumption, the index size and the query answering as well.

GBLENDER: Towards Blending Visual Query Formulation and Query Processing in Graph Databases

Changjiu Jin, Sourav S Bhowmick, Xiaokui Xiao, James Cheng (Nanyang Technological University), Byron Choi (Hong Kong Baptist University)

Given a graph database D and a query graph g, an exact subgraph matching query asks for the set S of graphs in D that contain g as a subgraph. This type of queries find important applications in several domains such as bioinformatics and chemoinformatics, where users are generally not familiar with complex graph query languages. Consequently, user-friendly visual interfaces that support query graph construction can reduce the burden of data retrieval for these users. Existing techniques for subgraph matching queries built on top of such visual framework are designed to optimize the time required in retrieving the result set S from D, assuming that the whole query graph has been constructed. This leads to sub-optimal system response time as the query processing is initiated only after the user has finished drawing the query graph. In this paper, we take the first step towards exploring a novel graph query processing paradigm, where instead of processing a query graph after its construction, it interleaves visual query construction and processing to improve system response time. To realize this, we present an algorithm called GBLENDER that prunes false results and prefetches partial query results by exploiting the latency offered by the visual query formulation. It employs a novel action-aware indexing scheme, which exploits users' interaction characteristics with visual interfaces to support efficient retrieval. Extensive experiments on both real and synthetic datasets demonstrate the effectiveness and efficiency of our solution.

Computing Label Constraint Reachability in Graph Databases

Ruoming Jin, Hui Hong (Kent State University), Haixun Wang (Microsoft Research, Asia), Yang Xiang, Ning Ruan (Kent State University)

Huge amount of graph data (biological networks, semantic web, ontologies, social networks) is being generated. Many of these real world graphs are edge-labeled graphs, i.e., each edge is associated with a label that denotes the relationship between the two vertices connected by the edge. Given this, a fundamental research problem on the labeled graphs is how to handle the label-constraint reachability query: if vertex u can reach vertex v through a path whose edge labels are constrained by a set of labels. In this work, we formally introduce the labeled reachability query problem, and provide two initial algorithms that categorize its computational complexity. We then generalize the transitive closure for the labeled graph, and introduce a new index framework which utilizes the directed spanning tree to compress the generalized transitive closure. We optimize the index for minimal memory cost and we propose a method which utilizes the directed maximal weighted spanning tree algorithm to maximally compress the generalized transitive closure. We establish an interesting link between the reachability problem and the geometric search problem, and we utilize the geometric search structures, such as multidimensional kd-trees or the range-search trees, to enable fast query processing. In our extensive experimental evaluation on both real and synthetic datasets, we find our tree-based index can significantly reduce the memory cost of the generalized transitive closure while still be able to answer the query very efficiently.

Pregel: A System for Large-Scale Graph Processing

Grzegorz Czajkowski, Greg Malewicz, Matthew Austern, Aart Bik, James Dehnert, Ilan Horn, Naty Leiser (Google)

Many practical computing problems concern large graphs. Standard examples include the Web graph and various social networks. The scale of these graphs -- in some cases billions of vertices, trillions of edges -- poses challenges to their efficient processing. In this paper we present a computational model suitable for this task. Programs are expressed as a sequence of iterations, in each of which a vertex can receive messages sent in the previous iteration, send messages to other vertices, and modify its own state and that of its outgoing edges. This vertex-centric approach is flexible enough to express a broad set of algorithms. The model has been designed for efficient, scalable and fault-tolerant implementation on clusters of thousands of commodity computers, and its implied synchronicity makes reasoning about programs easier. Distribution-related details are hidden behind an abstract API. The result is a framework for processing large graphs that is expressive and easy to program.

SIGMOD Research Session 4: Data Streams & Time-series Data

Chair: Alex Labrinidis (University of Pittsburgh)

Tuesday, 13:30 – 15:00

PR-Join: A Non-Blocking Join Achieving Higher Early Result Rate with Statistical Guarantees

Shimin Chen, Phillip Gibbons (Intel Labs Pittsburgh), Suman Nath (Microsoft Research)

Online aggregation is a promising solution to achieving fast early responses for interactive ad-hoc queries that compute aggregates on a large amount of data. Essential to the success of online aggregation is a good non-blocking join algorithm that enables both (i) high early result rates with statistical guarantees and (ii) fast end-to-end query times. We analyze existing non-blocking join algorithms and find that they all provide sub-optimal early result rates, and those with fast end-to-end times achieve them only by further sacrificing their early result rates. We propose a new non-blocking join algorithm, Partitioned expanding Ripple Join (PR-Join), which achieves a considerably higher early result rate than previous non-blocking joins, while also delivering fast end-to-end query times. PR-Join performs separate, ripple-like join operations on individual hash-based partitions, where the width of a ripple expands multiplicatively over time. This contrasts with the non-partitioned, fixed-width ripples of Block Ripple Join. Assuming, as in previous non-blocking join studies, that the input relations are in random order, PR-Join ensures representative early results that are amenable to statistical guarantees. We show both analytically and with real-machine experiments that PR-Join achieves over an order of magnitude higher early result rates than previous non-blocking joins. We also discuss the benefits of exploiting flash-based Solid State Drives (SSDs) as the temporary storage for non-blocking joins: Our experiments show that PR-Join achieves close to optimal end-to-end performance with SSDs as temporary storage. Finally, we extend PR-Join for joining finite data streams, and find that PR-Join achieves similar or higher result rates than RPJ, the state-of-the-art algorithm specialized for finite data streams.

PODS: A New Model and Processing Algorithms for Uncertain Data Streams

Thanh Tran, Liping Peng, Boduo Li, Yanlei Diao, Anna Liu (University of Massachusetts Amherst)

Uncertain data streams, where data is incomplete, imprecise, and even misleading, have been observed in many environments. Feeding such data streams to existing stream systems produces results of unknown quality, which is of paramount concern to monitoring applications. In this paper, we present the PODS system that supports stream processing for uncertain data naturally captured using continuous random variables. PODS employs a unique data model that is flexible and allows efficient computation. Built on this model, we develop evaluation techniques for complex relational operators, i.e., aggregates and joins, by exploring advanced statistical theory and approximation. Evaluation results show that our techniques can achieve high performance while satisfying accuracy requirements, and outperform a state-of-the-art sampling-based method significantly. A case study further shows that our techniques can enable a tornado detection system (for the first time) to produce detection results at stream speed and with much improved quality.

Fast Approximate Correlation for Massive Time-series Data

Abdullah Mueen (University of California, Riverside), Suman Nath, Jie Liu (Microsoft Research)

We consider the problem of computing all-pair correlations in a warehouse containing a large number (e.g., tens of thousands) of time-series data (or, signals). The problem arises in automatic discovery of patterns and anomalies in data intensive applications such as data center management, environmental monitoring, and scientific experiments. However, with existing techniques, solving the problem for a large stream warehouse is extremely expensive, due to the problem's inherent quadratic I/O and CPU complexities. We propose novel algorithms, based on Discrete Fourier Transform (DFT) and graph partitioning, to reduce the end-to-end response time of an all-pair correlation query. To minimize I/O cost, we partition a massive set of input signals into smaller batches such that caching the signals one batch at a time maximizes data reuse and minimizes disk I/O. To reduce CPU cost, we propose two approximation algorithms. Our first algorithm efficiently computes approximate correlation coefficients of similar signal pairs within a given error bound. The second algorithm efficiently identifies, without any false positives or negatives, all signal pairs with correlations above a given threshold. For many real applications, our approximate solutions are as useful as corresponding exact solutions, due to our strict error guarantees. However, compared to the state-of-the-art exact solution, our algorithms are up to 17 times faster for several real datasets.

An Algorithmic Approach to Event Summarization

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

Recently, much study has been directed toward summarizing event data, in the hope that the summary will lead us to a better understanding of the system that generates the events. However, instead of offering a global picture of the system, the summary obtained by most current approaches are piecewise, each describing an isolated snapshot of the system. We argue that the best summary, both in terms of its minimal description length and its interpretability, is the one obtained with the understanding of the internal dynamics of the system. Such understanding includes, for example, what are the internal states of the system, and how the system alternates among these states. In this paper, we adopt an algorithmic approach for event data summarization. More specifically, we use a hidden Markov model to describe the event generation process. We show that summarizing events based on the learned hidden Markov Model achieves short description length and high interpretability. Experiments show that our approach is both efficient and effective.

SIGMOD Research Session 5: Innovative Data Management

Chair: Mirek Riedewald (Northeastern University)

Tuesday, 13:30 – 15:00

Spreadsheet As a Relational Database Engine

Jerzy Tyszkiewicz (University of Warsaw)

Without any doubt, spreadsheets are the most commonly used applications for data management and analysis. Perhaps they are even among the most widely used computer applications of all kinds. However, the spreadsheet paradigm of computation still lacks sufficient theoretical analysis. In this paper we consider the relationship of spreadsheets to database systems. We demonstrate that a spreadsheet can play the role of a relational database engine, without any use of macros or built-in programming languages, merely by using spreadsheet formulas. We achieve that by implementing all operators of relational algebra by means of spreadsheet functions. Given a definition of a database (say in SQL), it is possible to construct a spreadsheet workbook with empty worksheets for data tables and worksheets filled with formulas for queries. Since then on, when the user enters, alters or deletes data in the data worksheets, the formulas in query worksheets automatically compute the actual results of the queries. Thus, the spreadsheet serves as data storage and executes SQL queries, and therefore acts as a relational database engine. Syntactically and semantically, the paper is based on Microsoft Excel (TM) 2003 version, because so far there is no formal model of spreadsheets that might be used for that purpose. However, the presented constructions work in other spreadsheet systems, too.

Scalable Architecture and Query Optimization for Transaction-time DBs with Evolving Schemas

Hyun Moon (NEC Labs), Carlo Curino (MIT), Carlo Zaniolo (UCLA)

The problem of archiving and querying the history of a database is made more complex by the fact that, along with the database content, the database schema also evolves with time. Indeed, archival quality can only be guaranteed by storing past database contents using the schema versions under which they were originally created. This causes major usability and scalability problems in supporting the preservation retrieval and querying databases that have hundreds of schema versions---and web information systems and scientific databases normally accumulate that many versions in just a few years. Our PostPrima system solves the usability issue by letting users write their queries against the current schema (or other schema versions) and then performing for the users the rewriting and execution of their queries on all appropriate schema versions. Furthermore, PostPrima achieves scalability by using (i) relational tables called H-tables that store the histories of independent attributes as separate tables, (ii) suitable temporal indexing and clustering techniques, and (iii) temporal query optimization techniques on H-tables. In particular, PostPrima uses a novel technique called \emph{CoalNesT} that minimizes the need for performing temporal coalescing on tuples that were fragmented by schema changes. Extensive experiments show that the performance and scalability thus achieved greatly exceeds those obtained by previous approaches. The PostPrima technology is easily deployed in current DBMS using a master/client architecture; moreover, it is robust and effective over a wide spectrum of archival solutions, ranging from the basic archive\&query features considered for new SQL standards, to the much richer XML/XQuery capabilities proposed by temporal database researchers.

Data Conflict Resolution Using Trust Relationships

Wolfgang Gatterbauer, Dan Suciu (University of Washington)

In massively collaborative projects such as scientific or community databases, users often agree or disagree on the content of individual data items. On the other hand, trust relationships often exist between users, allowing them to accept or reject other users' beliefs. As those trust relationships become complex, however, it becomes difficult to define and compute a consistent snapshot of the conflicting information. Previous solutions to a related problem, the update reconciliation problem, are dependent on the order in which the updates are processed and, therefore, do not guarantee a globally consistent snapshot. This paper proposes the first principled solution to the automatic conflict resolution problem in a community database. Our semantics is based on the certain tuples of all stable models of a logic program. While evaluating stable models in general is well known to be hard, even for very simple logic programs, we show that the conflict resolution problem admits a PTIME solution; to the best of our knowledge, ours is the first PTIME algorithm that allows conflict resolution in a principled way. We further discuss extensions to negative beliefs and prove that some of these extensions are hard.

Analyzing the Energy Efficiency of a Database Server

Dimitris Tsirogiannis (University of Toronto), Stavros Harizopoulos, Mehul Shah (HP Labs)

Rising energy costs in large data centers are driving an agenda for energy-efficient computing. In this paper, we focus on the role of database software in affecting, and, ultimately, improving the energy efficiency of a server. We first characterize the power-use profiles of database operators under different configuration parameters. We find that common database operations can exercise the full dynamic power range of a server, and that the CPU power consumption of certain operators, for the same CPU utilization, can differ by as much as 45%. We then experiment with several classes of database systems and storage managers, varying parameters that span from different query plans to compression algorithms, and from physical layout to CPU frequency and operating system scheduling. Contrary to what recent work in database systems may have suggested, we observe that within a single physical node, the most energy-efficient configuration is typically the highest performing one. We explain under which circumstances this is not the case and whether these circumstances warrant a re-engineering of existing systems. Our results reveal opportunities for

cross-node, energy-efficiency techniques but also inefficiencies in server-grade components that alternative or new hardware should address.

SIGMOD Research Session 6: Location & Sensor Based Data

Chair: Gottfried Vossen (WWU Munster)

Tuesday, 13:30 – 15:00

Processing Proximity Relations in Road Networks

Zhengdao Xu, Arno Jacobsen (University of Toronto)

Applications ranging from location-based services to multi-player online gaming require continuous query support to monitor, track, and detect events of interest among sets of moving objects. Examples are alerting capabilities for detecting whether the distance, the travel cost, or the travel time among a set of moving objects exceeds a threshold. These types of queries are driven by continuous streams of location updates, simultaneously evaluated over many queries. In this paper, we define three types of proximity relations that induce location constraints to model continuous spatio-temporal queries among sets of moving objects in road networks. Our focus lies on evaluating a large number of continuous queries simultaneously. We introduce a novel moving object indexing technique that together with a novel road network partitioning scheme restricts computations within the partial road network. These techniques reduce query processing overhead by more than 95%. Experiments over real-world data sets show that our approach is twenty times faster than a baseline algorithm.

Searching Trajectories by Locations - An Efficiency Study

Zaiben Chen, Heng Tao Shen, Xiaofang Zhou (The University of Queensland), Yu Zheng, Xing Xie (Microsoft Research Asia)

Trajectory search has long been an attractive and challenging topic which blooms various interesting applications in spatial-temporal databases. In this work, we study a new problem of searching trajectories by locations, in which context the query is only a small set of locations with or without an order specified, while the target is to find the k Best-Connected Trajectories (k-BCT) from a database such that the k-BCT best connect the designated locations geographically. Different from the conventional trajectory searches that look for similar trajectories w.r.t. shape or other criteria by using a sample query trajectory, we focus on the goodness of connection provided by a trajectory to the specified query locations. This new query can benefit users in many novel applications such as trip planning. In our work, we firstly define a new similarity function for measuring how well a trajectory connects the query locations, with both spatial distance and order constraint being considered. Upon the observation that the number of query locations is normally small (e.g. 10 or less) since it is impractical for a user to input too many locations, we analyze the feasibility of using a general-purpose spatial index to achieve efficient k-BCT searches, based on a simple Incremental k-NN based Algorithm (IKNN). The IKNN effectively prunes and refines trajectories by using the devised lower bound and upper bound of similarity. Our contributions mainly lie in adapting the best-first and depth-first k-NN algorithms to the basic IKNN properly, and more importantly ensuring the efficiency in both search effort and memory usage. An in-depth study on the adaption and its efficiency is provided. Further optimization is also presented to accelerate the IKNN algorithm. Finally, we verify the efficiency of the algorithm by extensive experiments.

Processing Continuous Join Queries in Sensor Networks: a Filtering Approach

Mirco Stern, Erik Buchmann, Klemens Bohm (Universitaet Karlsruhe)

While join processing in wireless sensor networks has received a lot of attention recently, current solutions do not work well for continuous queries. In those networks however, continuous queries are the rule. To minimize the communication costs of join processing, it is important to not ship non-joining tuples. In order to know which tuples do not join, prior work has proposed a precomputation step. For continuous queries however, repeating the precomputation for each execution is unnecessary and leaves aside that data tends to be temporally correlated. In this paper, we present a

filtering approach for the processing of continuous join queries. We propose to keep the filters and to maintain them. The main problems are determining the sizes of the filters and deciding which filters to update. Simplistic approaches result in bad performance. We show how to compute solutions that are optimal. Experiments on real-world sensor data indicate that our method performs close to a theoretical optimum and consistently outperforms state-of-the-art join approaches.

TACO: Tunable Approximate Computation of Outliers in wireless sensor networks

Nikos Giatrakos (University of Piraeus), Yannis Kotidis (Athens University of Economics and Business), Antonios Deligiannakis (Technical University of Crete), Vasilis Vassalos (Athens University of Economics and Business), Yannis Theodoridis (University of Piraeus)

Wireless sensor networks are becoming increasingly popular for a variety of applications. Users are frequently faced with the surprising discovery that readings produced by the sensing elements of their motes are often contaminated with outliers. Outlier readings can severely affect applications that rely on timely and reliable sensory data in order to provide the desired functionality. As a consequence, there is a recent trend to explore how techniques that identify outlier values based on their similarity to other readings in the network can be applied to sensory data cleaning. Unfortunately, most of these approaches incur an overwhelming communication overhead, which limits their practicality. In this paper we introduce an in-network outlier detection framework, based on locality sensitive hashing, extended with a novel boosting process as well as efficient load balancing and comparison pruning mechanisms. Our method trades off bandwidth for accuracy in a straightforward manner and supports many intuitive similarity metrics. Our experiments demonstrate that our framework can reliably identify outlier readings using a fraction of the bandwidth and energy that would otherwise be required.

SIGMOD Research Session 7: Probabilistic & Uncertain Data

Chair: Lei Chen (Hong Kong University of Science & Technology)

Tuesday, 15:30 – 17:00

GRN Model of Probabilistic Databases: Construction, Transition and Querying

Ruiwen Chen, Yongyi Mao, Iluju Kiringa (University of Ottawa)

Under the tuple-level uncertainty paradigm, we formalize the use of a novel graphical model, Generator-Recognizer Network (GRN), as a model of probabilistic databases. The GRN modeling framework extends the existing graphical models of probabilistic databases and is capable of representing a much wider range of dependency structures. In this framework, we show that a GRN representation of a database may undergo transitions induced by querying or constraining. We formalize the procedures of these two types of transitions so that the resulting graphical models after transitions remain as GRNs. This formalism makes GRN a self-contained modeling framework and a closed representation system for probabilistic databases -- a property that is lacking in most existing database models. In addition, we show that formally exploiting the transitional mechanisms allows a systematic approach to the construction of GRN models for arbitrary probabilistic databases at arbitrary stages. Advantages of GRNs in query processing are also demonstrated.

Consistent Query Answers in Inconsistent Probabilistic Databases

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

Efficient and effective manipulation of probabilistic data has become increasingly important recently due to many real applications that involve the data uncertainty. This is especially crucial when probabilistic data collected from different sources disagree with each other and incur inconsistency. This paper studies the problem of retrieving consistent answers to a series of query types in an inconsistent probabilistic databaseExtensive experiments have been conducted to demonstrate the efficiency and effectiveness of our proposed approaches.

Threshold Query Optimization for Uncertain Data

Yinian Qi, Rohit Jain, Sarvjeet Singh, Sunil Prabhakar (Purdue University)

The probabilistic threshold query (PTQ) is one of the most common queries in uncertain databases, where all results satisfying the query with probabilities that meet the threshold requirement are returned. PTQ is used widely in nearest-neighbor queries, range queries, ranking queries, etc. In this

paper, we investigate the general PTQ for arbitrary SQL queries that involve selections, projections and joins. Our uncertain database model is one that combines both attribute and tuple uncertainty as well as correlations between arbitrary attribute sets. We address the PTQ optimization problem that aims at improving the efficiency of PTQ query execution by pruning tuples that cannot be PTQ results as early as possible. We propose general optimization rules as well as rules specifically for selections, projections and joins. We introduce a threshold operator (τ-operator) to the query plan and argue it is generally most desirable to push down the τ-operator as much as possible. Our PTQ optimizations are evaluated in a real uncertain database management system. Our experiments on both real and synthetic data sets show that the optimizations improve the PTQ query processing time.

Probabilistic String Similarity Joins

Jeffrey Jestes, Feifei Li (Florida State University) Zhepeng Yan, Ke Yi (HKUST)

Edit distance based string similarity join is a fundamental operator in string databases. Increasingly, many applications in data cleaning, data integration, and scientific computing have to deal with fuzzy information in string attributes. Despite the intensive efforts devoted in processing (deterministic) string joins and managing probabilistic data respectively, modeling and processing probabilistic strings is still a largely unexplored territory. This work studies the string join problem in probabilistic string databases, using the expected edit distance (EED) as the similarity measure. We first discuss two probabilistic string models to capture the fuzziness in string values in real-world applications. The string-level model is complete, but may be expensive to represent and process. The character-level model has a much more succinct representation when uncertainty in strings only exists at certain positions. Since computing the EED between two probabilistic strings is prohibitively expensive, we have designed efficient and effective pruning techniques that can be easily implemented in existing relational database engines for both models. Extensive experiments on real data have demonstrated order of-magnitude improvements of our approaches over the baseline.

SIGMOD Research Session 8: Leveraging Hardware for Data management

Chair: Anastassia Ailamaki (EPFL, Zurich)

Tuesday, 15:30 – 17:00

FAST: Fast Architecture Sensitive Tree Search on Modern CPUs and GPUs

Changkyu Kim, Jatin Chhugani, Nadathur Satish (Intel), Eric Sedlar (Oracle) Anthony Nguyen (Intel) Tim Kaldewey (Oracle), Victor Lee (Intel), Scott Brandt (University of California, Santa Cruz), Pradeep Dubey (Intel)

In-memory tree structured index search is a fundamental database operation. Modern processors provide tremendous computing power by integrating multiple cores, each with wide vector (SIMD) unit. There has been much work to exploit modern processor architectures for database primitives like scan, sort, join and aggregation. However, unlike other primitives, index tree search presents significant challenges due to irregular and unpredictable data accesses in tree traversal. In this paper, we present FAST -- an extremely fast architecture sensitive layout of the index tree. FAST is a binary tree logically organized to optimize for architecture features like page size, cache line size, and SIMD width of the underlying hardware. FAST eliminates the impact of memory latency, and exploit thread-level and data-level parallelism on both CPUs and GPUs to achieve 50 million (CPU) and 85 million (GPU) queries per second, 5X (CPU) and 1.7X (GPU) better than the best previously reported performance on the same architectures. FAST also supports efficient bulk updates by rebuilding index trees in less than 0.1 seconds for datasets as large as 64M keys and naturally integrates compression techniques, overcoming the memory bandwidth bottleneck and achieving a 6X performance improvement over uncompressed index search for large keys on CPUs.

Fast In-Memory Sort on Modern CPUs and GPUs: A Case for Bandwidth-Oblivious SIMD Sort

Nadathur Satish, Changkyu Kim, Jatin Chhugani, Anthony Nguyen, Victor Lee, Daehyun Kim, Pradeep Dubey (Intel)

Sort is a fundamental kernel used in many database operations. In-memory sorts are now feasible; sort performance is limited by compute flops and main memory bandwidth rather than I/O. In this

paper, we present a competitive analysis of comparison and non-comparison based sorting algorithms on two modern architectures - the latest CPU and GPU architectures. We propose novel CPU radix sort and GPU merge sort implementations which are 2X faster than previously published results. We perform a fair comparison of the algorithms using these best performing implementations on both architectures. While radix sort is faster on current architectures, the gap narrows from CPU to GPU architectures. Merge sort performs better than radix sort on keys larger than 9 bytes on CPUs and 8 bytes on GPUs. We present analytical models for analyzing the performance of our implementations in terms of how friendly they are to architectural features such as core count, SIMD and bandwidth. We show that our obtained performance results are successfully predicted by our models. Our analysis points to merge sort winning over radix sort on future architectures due to its efficient utilization of SIMD parallelism and low bandwidth utilization. We simulate a 64-core platform with varying SIMD widths under constant bandwidth per core constraints, and show that even on data sizes as large as one trillion records, merge sort performance on 8-byte keys is 1.5X better than radix sort on 2048-bit wide SIMD, and the performance gap increases further to 3X for 16-byte keys. Therefore, merge sort should be the sorting method of choice for future databases.

Page-Differential Logging: An Efficient and DBMS-independent Approach for Storing Data into Flash Memory

Yi-Reun Kim, Kyu-Young Whang (KAIST), Il-Yeol Song (Drexel University)

Flash memory is widely used as the secondary storage in lightweight computing devices due to its outstanding advantages over magnetic disks. Flash memory has many access characteristics different from those of magnetic disks, and how to take advantage of them is becoming an important research issue. There are two existing approaches to storing data into flash memory: page-based and log-based. The former has good performance for read operations, but poor performance for write operations. In contrast, the latter has good performance for write operations when updates are light, but poor performance for read operations. In this paper, we propose a new method of storing data, called {\it page-differential logging}, for flash-based storage systems that solves the drawbacks of the two methods. The primary characteristics of our method are: (1) writing only the difference (which we define as the {\it page-differential}) between the original page in flash memory and the up-to-date page in memory; (2) computing and writing the page-differential only once at the time the page needs to be reflected into flash memory. The former contrasts with existing page-based methods that write the whole page including both changed and unchanged parts of data or from log-based ones that keep track of the history of all the changes in a page. Our method allows existing disk-based DBMSs to be reused as flash-based DBMSs just by modifying the flash memory driver, i.e., it is DBMS-independent. Experimental results show that the proposed method is superior in I/O performance, except for some special cases, to existing ones. Specifically, it improves the performance of various mixes of read-only and update operations by 0.6 (the special case when all transactions are read-only on updated pages) ~ 1.7 times over the page-based method and by 1.4 ~ 2.5 times over the log-based one for synthetic data of approximately 1 Gbytes.

Similarity Search and Locality Sensitive Hashing using Ternary Content Addressable Memories

Rajendra Shinde, Ashish Goel (Stanford University), Pankaj Gupta (Twitter Inc.), Debojyoti Dutta (Cisco Systems Inc.)

Similarity search is a widely used kernel in various data mining and machine learning applications including those from computational biology, web search/clustering. Nearest neighbor (NN) methods are often used to retrieve similar entries, given a query. While there exist efficient techniques for exact query lookup, the best known similarity search (including NN search) techniques offer little improvement over brute force techniques in large dimensional spaces.Fast solutions to the approximate NN problem include Locality Sensitive Hashing (LSH) based techniques, which need storage O(n^{1+\rho}), and query time O(n^{\rho}) where n is the size of the database and \rho > 0. In this work we present a new technique of solving the approximate NN problem in Euclidean space using a Ternary Content Addressable Memory (TCAM), in near linear space and O(1) query time. TCAMs are high performance associative memories widely used in networking applications such as address lookups and access control lists. It can query for a bit vector within a database of ternary

vectors, where every bit position represents 0, 1 or *, a wild card representing either a 0 or a 1. We design a new variant of LSH for TCAMs, called Ternary Locality Sensitive Hashing (TLSH), by hashing database entries represented by vectors in the Euclidean space into {0,1,*}, to solve an instance of the approximate NN problem with 1 TCAM access. We validate our claims with extensive simulations using both real world (Wikipedia) as well as synthetic (but illustrative) datasets. We observe that a TCAM of width 288 bits solves the approximate NN problem on a database of size 1M points with few false positives and false negatives. Finally, we design an experiment with TCAMs within a Cisco ethernet switch to validate that indeed TLSH could be used to perform 1.5M queries per second per 1Gbps port.

SIGMOD Research Session 9: Web Data Integration

Chair: Magdalena Balazinska (University of Washington)

Wednesday, 8:30 – 10:00

Automatically Incorporating New Sources in Keyword Search-Based Data Integration

Partha Talukdar, Zachary Ives (University of Pennsylvania), Fernando Pereira (Google)

One of the more important use cases of data integration today is across scientific communities, e.g., in astronomy, medicine, and biology. A major challenge in all of these fields is that they are constantly evolving: existing data is annotated or revised, new data is contributed, and new databases are made available on the Web. From the perspective of the user, this can be a major headache: the data they want may initially be spread across many databases, and such data needs to be integrated. Worse, even if they are given a solution that integrates the current state of the source databases (e.g., by using virtual integration techniques), new data sources may emerge, with new data items the user would like to incorporate. In this paper we build upon recent ideas for creating integrated views over data sources using keyword search techniques, ranked answers, and user feedback; and consider how to automatically discover when a new data source has content relevant to a user's view --- in essence, performing automatic data integration in this setting. Our architecture can accommodate a variety of methods to discover related attributes, including algorithms from the machine learning community as well as off-the-shelf schema matchers. The user may provide feedback on the suggested new results, helping the system repair any bad alignments or drop the relevance score for new sources that are not useful. We evaluate our approach using state-of-the-art schema matching tools as components, over real data and schemas from the bioinformatics domain. We additionally discuss how our architecture can be adapted to more traditional settings with a mediated schema.

Active Knowledge: Dynamically Enriching RDF Knowledge Bases by Web Services

Nicoleta Preda (Max-Planck Institute), Gjergji Kasneci, Fabian Suchanek (Microsoft Research), Thomas Neumann (Max-Planck Institute), Wenjun Yuan (University of Hong Kong), Gerhard Weikum (Max-Planck Institute)

The proliferation of knowledge-sharing communities and the advances in information extraction have enabled the construction of large knowledge bases using the RDF data model to represent entities and relationships. However, as the Web and its latently embedded facts evolve, a knowledge base can never be complete and up-to-date. On the other hand, a rapidly increasing suite of Web services provide access to timely and high-quality information, but this is encapsulated by the service interface. We propose to leverage the information that could be dynamically obtained from Web services in order to enrich RDF knowledge bases on the fly whenever the knowledge base does not suffice to answer a user query. To this end, we develop a sound framework for appropriately generating queries to encapsulatedWeb services and efficient algorithms for query execution and result integration. The query generator composes sequences of function calls based on the available service interfaces. As Web service calls are expensive, our method aims to minimize the number of calls in order to retrieve results with sufficient recall. Our approach is fully implemented in a complete prototype system. The user can query and browse the RDF knowledge base as if it already contained all the facts from the Web services. This data, however, is gathered and integrated on the fly, transparently to the user. We demonstrate the viability and efficiency of our approach in experiments based on real-life data provided by popular Web services.

Schema Clustering and Retrieval for Multi-domain Pay-As-You-Go Data Integration Systems

Hatem Mahmoud, Ashraf Aboulnaga (University of Waterloo)

A data integration system offers a single interface to multiple data sources. Many application contexts (e.g., searching structured data on the web) involve the integration of large numbers of data sources. On the web scale, it is impractical to use manual or semi-automatic integration methods, therefore a pay-as-you-go approach is more appropriate on that scale. In a pay-as-you-go approach, fully-automatic approximate integration bootstraps the system, then further enhancements can be later introduced if required. Previous research already investigated fully-automatic approximate integration, and probabilistic models were proposed to deal with uncertainty in it. However, existing techniques of data integration work on schemas that belong to the same domain. On the web scale, it is impractical to classify schemas manually or using semi-automatic techniques. In this paper, we propose techniques for clustering schemas, and we deal with uncertainty in clustering via a probabilistic model. We also propose a query processor based on naive Bayes classification to determine, given a keyword query, the most relevant domains to that query. We show how to integrate our probabilistic model with related work on data integration with uncertainty. Finally, we evaluate our system on hundreds of extremely heterogeneous and overlapping data sources extracted from the web, and we show that our approach works effectively with different types of schemas, domains, and queries.

Expressive and Flexible Access to Web-Extracted Data: A Keyword-based Structured Query Language

Jeffrey Pound, Ihab Ilyas, Grant Weddell (University of Waterloo)

The automated extraction of structured data from Web sources often leads to large heterogeneous knowledge bases (KB), with data and schema items numbering in the hundreds of thousands or millions. Formulating information needs with conventional structured query languages is difficult due to the sheer size of schema information available to the user. We address this challenge by proposing a new query language that blends keyword search with structured query processing over large schema graphs with rich semantics. We introduce a new formalism for structured queries based on keywords that combines the flexibility of keyword search with the expressiveness of structures queries. We propose a solution to the resulting disambiguation problem caused by introducing keywords as primitives in a structured query language. We show how expressions in our proposed language can be rewritten using the vocabulary of the web-extracted KB, and how different possible rewritings can be ranked based on their syntactic relationship to the keywords in the query as well as their semantic relationships in the underlying KB. An extensive experimental study demonstrates the efficiency and effectiveness of our approach. Additionally, we show how our query language fits into QUICK, an end-to-end information system that integrates web-extracted data graphs with full-text search. In this system, the rewritten query describes an arbitrary topic of interest for which corresponding entities, and documents relevant to the entities, are efficiently retrieved.

SIGMOD Research Session 10: Social Networks & Community Data

Chair: Susan Davidson (University of Pennsylvania)

Wednesday, 8:30 – 10:00

Multiple Features Fusion for Social Media Applications

Bin Cui (Peking University), Anthony Tung (National University of Singapore), Ce Zhang, Zhe Zhao (Peking University)

With the continued growth of multimedia data on the Internet, multimedia data management is attracting increasing attention. More recently, the emergence of social media as a crucial paradigm has posed new challenges to the research and industry communities, where media are designed to be disseminated through social interaction. Recent literature has noted the generality of multiple features in the social media environment, such as textual, visual and user information. However, most of the studies employ only a relatively simple mechanism to merge the features rather than fully exploit feature correlation for social media applications. In this paper, we propose a novel

approach to fusing multiple features and their correlations for similarity evaluation. Specifically, we first build a Feature Interaction Graph (FIG) by taking features as nodes and the correlations between them as edges. Then, we employ a probabilistic model based on Markov Random Field to describe the graph for similarity measure between multimedia objects. Using that, we design an efficient retrieval algorithm for large social media data. Further, we integrate temporal information into the probabilistic model for social media recommendation. We evaluate our approach using a large real-life corpus collected from Flickr, and the experimental results indicate the superiority of our proposed method over state-of-the-art techniques.

Finding Maximal Cliques in Massive Networks by H*-Graph

James Cheng, Yiping Ke, Ada Fu, Jeffrey Xu Yu (Chinese University of Hong Kong), Linhong Zhu (NTU, Singapore)

Maximal clique enumeration (MCE) is a fundamental problem in graph theory and has important applications in many areas such as social network analysis and bioinformatics. The problem is extensively studied. However, the best existing algorithms require memory space linear in the size of the input graph. This has become a serious concern in view of the massive volume of today's fast-growing networks. Since MCE requires random accesses to different parts of a large graph, it is difficult to divide the graph into smaller parts and process one part at a time, because either the result may be incorrect and incomplete, or it incurs huge cost on merging the results from different parts. We propose the notion of H*-graph, which defines the core of a network and also extends to encompass the neighborhood of the core for MCE. We propose a novel recursive MCE algorithm (RMCE), which uses the H*-graph to bound the memory usage as well as to guide the process of the recursive steps. We show that the size of the H*-graph is small for real-world networks. We then prove both the correctness and the completeness of the result computed by RMCE. Our extensive experimental studies verify that RMCE efficiently processes large real-world networks that are infeasible for existing algorithms to compute. We also show that the H*-graph captures important properties of the network; thus, updating the maximal cliques in the H*-graph retains the most essential information, with a low update cost, when it is infeasible to perform update on the entire network.

K-Isomorphism: Privacy Preservation in Network Publication against structural attack

James Cheng, Ada Fu, Jia Liu (Chinese University of Hong Kong)

Privacy protection in social network publication has been considered an important problem. A lot of real life social network information is available in the public, and protecting the privacy of the individuals involved is necessary. We consider the problem of structural attack where the adversary has the knowledge of a neighborhood graph around the vertices of targeted individuals. Previous studies found that structural attack can be highly successful on identifying the vertex of an individual when the network or graph data is released with only the names and other identification information removed. We propose a novel method to safeguard the node and link information of the network. Our anonymization method involves both edge addition and deletion. Empirical studies show that the utility of the anonymized data is maintained, while the privacy is guaranteed at our desired level.

Load-Balanced Query Dissemination in Democratic Communities

Emiran Curtmola, Alin Deutsch (UCSD), K.K. Ramakrishnan, Divesh Srivastava (AT&T Labs – Research)

As the web evolves, it is becoming easier to form communities based on shared interests, and to create and publish data on a wide variety of topics. With this democratization of information creation comes the natural desire to make one's data accessible for querying within the community and also be able to query the global collection that is the union of all local data collections of others within the community. In order to fully deliver on the promise of free data exchange and prevent user censorship or discrimination, any community-supporting infrastructure needs to enforce the key requirement to preserve privacy of the association of content providers with potential sensitive published information. This privacy preserving publishing requirement precludes some obvious approaches that reuse and build on existing centralized technologies. We propose a novel privacy-

preserving enabling distributed infrastructure in which data resides only with the publishers owning it. The infrastructure disseminates user queries to publishers, who answer them at their own discretion. The infrastructure enforces a publisher k-anonymity guarantee, which prevents leakage of information about which publishers are capable of answering a certain query. Given the virtual nature of the global data collection, we study the challenging problem of efficiently locating publishers in the community that contain data items matching a specified query. We propose a distributed index structure that is organized as a union of Query Dissemination Trees (QDTs), and realized on a network overlay. Each QDT has data publishers as its leaf nodes, and overlay network nodes as its internal nodes; each internal node routes queries to publishers, based on a summary of the data advertised by publishers in its subtrees. We experimentally evaluate design tradeoffs, and demonstrate that we can maximize throughput by preventing any overlay network node from becoming a bottleneck.

SIGMOD Research Session 11: Scalable Data Analytics

Chair: Chris Olston (Yahoo! Research)
Wednesday, 8:30 – 10:00

Automatic Contention Detection and Amelioration for Data-Intensive Operations

John Cieslewicz, Kenneth Ross, Kyoho Satsumi, Yang Ye (Columbia University)

To take full advantage of the parallelism offered by a multi-core machine, one must write parallel code. Writing parallel code is difficult. Even when one writes correct code, there are numerous performance pitfalls. For example, an unrecognized data hotspot could mean that all threads effectively serialize their access to the hotspot, and throughput is dramatically reduced. Previous work has demonstrated that database operations suffer from such hotspots when naively implemented to run in parallel on a multi-core processor. In this paper, we aim to provide a generic framework for performing certain kinds of concurrent database operations in parallel. The formalism is similar to user-defined aggregates and Google's MapReduce in that users specify certain functions for parts of the computation that need to be performed over large volumes of data. We provide infrastructure that allows multiple threads on a multi-core machine to concurrently perform read and write operations on shared data structures, automatically mitigating hotspots and other performance hazards. Our goal is not to squeeze the last drop of performance out of a particular platform. Rather, we aim to provide a framework within which a programmer can, without detailed knowledge of concurrent and parallel programming, develop code that efficiently utilizes a multi-core machine.

Efficient Parallel Set-Similarity Joins Using MapReduce

Rares Vernica, Michael Carey (Univ of California, Irvine), Chen Li (University of California, Irvine and BiMaple)

In this paper we study how to efficiently perform set-similarity joins in parallel using the popular MapReduce framework. We propose a 3-stage approach for end-to-end set-similarity joins. We take as input a set of records and output a set of joined records based on a set-similarity condition. We efficiently partition the data across nodes in order to balance the workload and minimize the need for replication. We study both self-join and R-S join cases, and show how to carefully control the amount of data kept in main memory on each node. We also propose solutions for the case where, even if we use the most fine-grained partitioning, the data still does not fit in the main memory of a node. We report results from extensive experiments on real datasets, synthetically increased in size, to evaluate the speedup and scaleup properties of the proposed algorithms using Hadoop.

ParaTimer: A Progress Indicator for MapReduce DAGs

Kristi Morton, Magdalena Balazinska, Dan Grossman (University of Washington)

Accurate progress estimation for parallel queries is a challenging problem that has received only limited attention. The challenges are especially great when users are interested in estimates of time remaining rather than a count of processed records. Previous work has focused only on time-

remaining estimates for single-site queries and a very limited class of parallel queries, which exclude joins, data skew, node failures, and other important challenges that arise in practice. In this paper, we present ParaTimer, a comprehensive time-remaining indicator for parallel queries. ParaTimer builds on previous techniques and makes two key contributions. First, it estimates the progress of parallel queries that include joins, which requires a radically different approach than in prior work. Second, it handles a variety of real systems challenges such as failures and data skew. To handle unexpected changes in query execution times due to runtime condition changes, ParaTimer provides users not only with one but with a set of time-remaining estimates, each one corresponding to a different carefully selected scenario. Several parallel data processing systems exist. In this paper, we target environments where declarative Pig Latin queries are translated into MapReduce DAGs. We implement our estimator in the Pig system and demonstrate its performance on experiments running on a real, small-scale cluster.

The DataPath System: A Data-Centric Analytic Processing Engine for Large Data Warehouses

Subi Arumugam, Alin Dobra (University of Florida), Christopher Jermaine, Luis Perez, Niketan Pansare (Rice University)

Since the 1970's, database systems have always been ``compute-centric". That is, when a computation needs data, it requests it, and the data are pulled through the system. We believe that pull-based database system design necessarily leaves considerable performance on the table, for two reasons. First, requests for data naturally incur high latency as the data are pulled through the memory hierarchy, and second, it makes it difficult or impossible for multiple queries or operations that are interested in the same data to amortize the bandwidth and latency costs associated with their data access. In this paper, we describe a purely-push based, research prototype database system called DataPath. DataPath is ``data-centric". In DataPath, queries do not request data. Instead, data are automatically pushed onto processors, where they are then processed by any interested computation. We show experimentally on a multi-terabyte benchmark that this basic design principle makes for a very lean and fast database system, with excellent performance on inexpensive hardware. In particular, it results in an architecture that is able to efficiently process many queries concurrently.

SIGMOD Research Session 12: Advanced Query Processing

Chair: Jiaheng Lu (Renmin University)

Wednesday, 10:30 – 12:00

Variance Aware Optimization of Parameterized Queries

Surajit Chaudhuri (Microsoft Research), Hongrae Lee (University of British Columbia), Vivek Narasayya (Microsoft Research)

Parameterized queries are commonly used in database applications. In a parameterized query, the same SQL statement is potentially executed multiple times with different parameter values. In today's DBMSs the query optimizer typically chooses an execution plan that is reused for multiple instances of the same query. A key problem is that even if a plan with low average cost across instances is chosen, its variance can be high, which can lead to unpredictable performance. In this paper, we describe techniques for selecting a plan that can better address the trade-off between the average and variance cost across instances of a parameterized query. We show how to efficiently compute the skyline of non-dominated plans in the average-variance space. We have implemented our techniques on top of a commercial DBMS. Our experimental results on benchmark and real-world decision support queries show the effectiveness of our techniques.

Positional Update Handling in Column Stores

Sandor Heman, Marcin Zukowski (VectorWise), Niels Nes, Lefteris Sidirourgos, Peter Boncz (CWI)

In this paper we investigate techniques that allow for on-line updates to columnar databases, leaving intact their high read-only performance. Rather than keeping differential structures organized by the table key values, the core proposition of this paper is that this can better be done by keeping track of the tuple position of the modifications. Not only does this minimize the computational overhead of merging in differences into read-only queries, but this makes the differential structure oblivious of

the value of the order keys, giving the advantage of smaller memory consumption, and – critically in a column store – the ability to avoid disk I/O for retrieving the order keys in read-only queries that otherwise do not need them. We describe a new data structure for maintaining such positional updates, called the Positional Delta Tree (PDT), and describe detailed algorithms for PDT/column merging, updating PDTs and using PDTs in transaction management. In experiments with a fully functional columnar DBMS, we both perform microbenchmarks on PDTs, as well as show in a TPC-H workload that PDTs allow quick on-line updates, yet greatly reduce the impact of merging updates in read-only queries when comparing with classical value-based differential methods.

Durable Top-k Search in Document Archives

Leong Hou U, Nikos Mamoulis (University of Hong Kong), Klaus Berberich, Srikanta Bedathur (MPII)

We propose and study a new ranking problem in versioned databases. Consider a database of versioned objects which have different valid instances along a history (e.g., documents in an web archive). Durable top-k search finds the set of objects that are consistently in the top-k results of a query (e.g., a keyword query) throughout a given time interval (e.g., from June 2008 to May 2009). Existing work on temporal top-k queries mainly focuses on finding the most representative top-k elements within a time interval. As we discuss in the paper, such methods are not readily applicable to durable top-k queries. To address this need, we propose two different techniques that compute the durable top-k result. The first one is adapted from the classic top-k rank aggregation algorithm NRA. The second technique is based on a shared execution paradigm and is more efficient than the first approach. In addition, we also study an appropriate indexing technique for archived data. In our experiments, we use real data to demonstrate the efficiency of our solutions.

Ajax-based Report Pages as Incrementally Rendered Views

Yupeng FU (UCSD), Keith Kowalczykowski (App2you Inc), Yannis Papakonstantinou, Kevin Keliang Zhao, Kian Win Ong (UCSD)

Ajax-based programming enables web application performance and interface quality commensurate to that of desktop applications and far superior than that of purely server side applications. Its key advantage is Javascript-issued asynchronous server requests that lead to quick partial updates of the page and continuous uninterrupted user activity on the browser. While Ajax provides great advantages over the pure server side programming, it is more demanding and error prone as it is distributed programming between the browser's Javascript, which accesses the page state, and the server's Java, which accesses the database. Furthermore, realizing the partial update benefits requires the implementation of code that accounts for the effect of each action on the page. The problem is exacerbated when components (maps, calendars) are used. The SPRTF framework achieves the functionality of Ajax while keeping the simplicity of pure server side programming: It provides a database that encompasses the page state. Each request is processed using queries that have location-transparent access to the page state and the database. The new page state is the result of a nested query (suitably designed for the task), which with the help of a rendering layer turns into HTML and component state. Overall, the use of SPRTF leads to significantly less code. Such ease-of-development is due to SPRFT solving performance optimization problems that would otherwise need to be hand-coded by the developer: (1) How to efficiently mirror the page state to the server? (2) How to create the new page by incrementally updating the old page. Since the pages are fueled by queries we leverage years of database research on incremental view maintenance by creating optimization techniques appropriately extended for the needs of the problem (nesting, variability) and provide experimental and use case data showing that SPRTF achieves performance comparable to hand-coded Javascript/Java applications.

SIGMOD Research Session 13: Cloud Computing & Internet Scale Computing

Chair: Mehul Shah (HP Labs)

Wednesday, 10:30 – 12:00

An Evaluation of Alternative Architectures for Transaction Processing in the Cloud

Simon Loesing, Tim Kraska, Donald Kossmann (ETH Zurich)

Cloud computing promises a number of advantages for the deployment of data-intensive applications. One important promise is reduced cost with a pay-as-you-go business model. Another promise is scalability and elasticity by adding and removing nodes depending on the application requirements. This paper lists alternative architectures to effect cloud computing for database applications and reports on the results of a comprehensive evaluation of existing commercial cloud services that have adopted these architectures. The focus of this work is on transaction processing (i.e., read and update workloads), rather than analytics or OLAP workloads, which have recently gained a great deal of attention. The results are surprising in several ways. Most importantly, it seems that all major vendors have adopted a different architecture for their cloud services. As a result, the cost and performance of the services vary significantly depending on the workload.

Indexing Multi-dimensional Data in a Cloud System

Jinbao Wang, Hong Gao (Harbin Institute of Technology), Sai Wu, Beng chin Ooi (National University of Singapore)

Mapping mashups are becoming popular as Web 2.0 data representations where photos, blogs, video and other user generated data are associated with geo-position tags and are marked in a map such as Google and Yahoo Maps. These objects may also have multi-dimensional attributes that describe their properties. Many of such applications are typically hosted on Cloud platforms where data are partitioned based on some properties and are distributed over a network of commodity compute nodes that are connected using standard network hardward. Without efficient indexing structures as in centralized database systems, the Cloud based systems have to flood any search query to all compute nodes and scan the dataset stored in each node. Due to the huge number of compute nodes and the limited listener processes each node can support, it is not possible to maintain point-to-point persistent connection in the Cloud system. On the other hand, on-the-fly stablishment of connection between nodes incurs fairly high start-up costs and causes noticeable latency. In this paper, we propose a CAN based routing protocol and multi-dimensional R-tree based indexing scheme called the RT-CAN for supporting efficient routing and multi- dimensional query processing in a Cloud system. In our proposal, we make a simple assumption that each compute node uses an R-tree like indexing structure to index the data locally stored. We propose a query pattern conscious cost model that will select beneficial local R-tree nodes for publising over the RT-CAN. By keeping the number of persistently connected nodes small, and maintaining a global multi-dimensional search index, we can locate the compute nodes that may contain the answer with a few hops, making the scheme scalable in terms of data volume and number of compute nodes. Experiments on Amazon's EC2 show that our proposed routing protocol and indexing scheme are robust, efficient and scalable.

Low Overhead Concurrency Control in Partitioned DBMSs

Evan Jones (MIT), Daniel Abadi (Yale), Samuel Madden (MIT)

Database partitioning is a well known technique to improve the performance of distributed OLTP databases, since "single partition" transactions that access data on only one partition do not need distributed coordination with a protocol such as two-phase commit. For workloads that are amenable to partitioning, some argue that transactions should be executed serially on each partition without any concurrency at all. This strategy makes sense when data is stored in main memory and there are no disk or user stalls, since the CPU is well utilized and the substantial overhead of traditional concurrency control techniques such as two-phase locking can be avoided. Unfortunately, many OLTP applications have some transactions which span multiple partitions. This introduces network stalls in order to coordinate distributed transactions, which will limit the performance of a database that does not allow concurrency. In this paper, we compare two low overhead concurrency control schemes that relax the strict "no concurrency" requirement and allow partitions to work on other transactions during network stalls, yet have little to no cost in the common case when concurrency is not needed. The first is a light-weight locking scheme, and the second is an even lighter-weight type of speculative concurrency control that avoids the overhead of tracking reads and writes but sometimes performs work that eventually must be undone. We quantify the range of workloads over which each technique is beneficial, showing that speculative concurrency control generally outperforms locking as long as there are few aborts or few distributed transactions that involve multiple rounds of communication. On a modified TPC-C benchmark, speculative concurrency control can improve throughput relative to the other schemes by up to a factor of two.

Efficient Querying and Maintenance of Network Provenance at Internet-Scale

Wenchao Zhou, Micah Sherr, Tao Tao, Xiaozhou Li, Boon Thau Loo, Yun Mao (University of Pennsylvania)

Network accountability, forensic analysis, and failure diagnosis are becoming increasingly important for network management and security. Such capabilities often utilize {\em network provenance} -- the ability to issue queries over network meta-data. For example, network provenance may be used to trace the path a message traverses on the network as well as to determine how message data was derived and which parties were involved in its derivation. This paper presents the design and implementation of ExSPAN, a {\em generic and extensible} framework to achieve efficient network provenance in a distributed environment. We utilize the database notion of {\em data provenance} to ``explain" the existence of {\em any} network state, providing a versatile mechanism for network provenance. To achieve such flexibility, ExSPAN use declarative networking in which network protocols can be modeled as continuous queries over distributed streams and specified concisely in a declarative query language. We extend existing data models for provenance developed in database literature, to enable distribution at Internet-scale, and investigate numerous optimization techniques to maintain and query distributed network provenance efficiently. The ExSPAN prototype is developed using RapidNet, a declarative networking platform based on the emerging ns-3 toolkit. Our experiments over a simulated network and an actual deployment on a testbed environment demonstrate that our system can support a wide range of distributed provenance computations efficiently, resulting in significant reduction in bandwidth utilization compared to centralized approaches.

SIGMOD Research Session 14: Data Summarization

Chair: Yannis Papakonstantinou (UC San Diego)

Wednesday, 10:30 – 12:00

Hierarchically Organized Skew-Tolerant Histograms for Geographic Data Objects

Yohan Roh (SAIT, Samsung Electronics), Jae Ho Kim (KAIST), Yon Dohn Chung (Korea University), Jin Hyun Son (Hanyang University), Myoung Ho Kim (KAIST)

Histograms have been widely used for fast estimation of query result sizes in query optimization. In this paper, we propose a new histogram method, called the Skew-Tolerant Histogram (STHistogram) for two or three dimensional geographic data objects that are used in many real-world applications in practice. The proposed method provides a significantly enhanced accuracy in a robust manner even for the data set that has a highly skewed distribution. Our method detects hotspots present in various parts of a data set and exploits them in organizing histogram buckets. For this purpose, we first define the concept of a hotspot, and provide an algorithm that efficiently extracts hotspots from the given data set. Then, we present our histogram construction method that utilizes hotspot information. We also describe how to estimate query result sizes by using the proposed histogram. We show through extensive performance experiments that the proposed method provides better performance than other existing methods.

Logging Every Footstep: Quantile Summaries for the Entire History

Yufei Tao (Chinese University of Hong Kong), Ke Yi (HKUST), Sheng Cheng (Chinese University of Hong Kong), Jian Pei (Simon Fraser University), Feifei Li (Florida State University)

Quantiles are a crucial type of order statistics in databases. Extensive research has been focused on maintaining a space-efficient structure that can be used for approximate quantile computation as the underlying dataset is updated. The existing solutions, however, are designed to support only the current, most-updated, snapshot of the dataset. Queries on the past versions of the data cannot be answered. This paper studies the problem of {\em historical quantile search}. The objective is to enable $\eps$-approximate quantile retrieval on {\em any} snapshot of the dataset in history. The problem is very important in analyzing the evolution of a distribution, monitoring the quality of service, query optimization in temporal databases, and so on. We present the first formal results in the literature. First, we prove a novel theoretical lower bound on the space cost of supporting $\eps$-

approximate historical quantile queries. The bound reveals the fundamental difference between answering quantile queries about the past and those in a single snapshot. Second, we propose a structure for finding $\eps$-approximate historical quantiles, and show that it consumes more space than the lower bound by only a square-logarithmic factor. Extensive experiments demonstrate that in practice our technique performs much better than predicted by theory. In particular, the quantiles it returns are remarkably more accurate than the theoretical precision guarantee.

Continuous Sampling for Online Aggregation Over Multiple Queries
Sai Wu, Beng Chin Ooi, Kian-Lee Tan (National University of Singapore)

In this paper, we propose an online aggregation system called COSMOS (Continuous Sampling for Multiple queries in an Online aggregation System), to process multiple aggregate queries efficiently. In COSMOS, a dataset is first scrambled so that sequentially scanning the dataset gives rise to a stream of random samples for for all queries. Moreover, COSMOS organizes queries into a dissemination graph to exploit the dependencies across queries. In this way, aggregates of queries closer to the root (source of data flow) can potentially be used to compute the aggregates of descendent/dependent queries. COSMOS applies some statistical approach to combine answers from ancestor nodes to generate the online aggregates for a node. COSMOS also offers a partitioning strategy to further salvage intermediate answers. We have conducted an extensive experimental study using PostgreSQL on the TPC-H benchmark. Our results show the efficiency and effectiveness of COSMOS.

Histograms Reloaded: The Merits of Bucket Diversity
Carl-Christian Kanne, Guido Moerkotte (University of Mannheim)

Virtually all histograms store for each bucket the number of distinct values it contains and their average frequency. In this paper, we question this paradigm. We start out by investigating the estimation precision of three commercial database systems which also follow the above paradigm. It turns out that huge errors are quite common. We then intro- duce new bucket types and investigate their accuracy when building optimal histograms with them. The results are ambiguous. There is no clear winner among the bucket types. At this point, we (1) switch to heterogeneous histograms, where different buckets of the same histogram possibly are of different types, and (2) design more bucket types. The nice consequence of introducing heterogeneous histograms is that we can guarantee decent upper error bounds while at the same time heterogeneous histograms require far less space than homogeneous histograms.

SIGMOD Research Session 15: Probabilistic Data, Fuzzy Data, & Data Provenance

Chair: Martin Theobald (Max-Planck-Institut fur Informatik)
Thursday, 10:30 – 12:00

Lineage Processing over Correlated Probabilistic Databases
Bhargav Kanagal, Amol Deshpande (University of Maryland)

In this paper, we present a system for processing conjunctive queries over correlated tuple uncertain probabilistic databases. As with previous work, our query processor first determines the lineages of the output tuples and subsequently computes the probability of the lineage formulae. Unlike previous work however, we address the case when the input tuples are highly correlated. These correlations are captured via a forest of junction trees over the random variables corresponding to the uncertain tuples. We show that processing even read-once (tree structured) lineages (e.g., those generated by hierarchical conjunctive queries), which is polynomially computable over tuple independent probabilistic databases, is #P-complete for lightly correlated probabilistic databases e.g., Markov sequences. Thereby, exactly computing the probabilities is often infeasible. Despite this handicap, we develop algorithms for processing a lineage formula based on message passing over junction trees and characterize their complexity in terms of lwidth, a new parameter we introduce, that depends on the input lineage and the structure of the correlations in the data. For lineages that induce low lwidth, we compute exact probabilities and for those that induce large lwidths, we resort to

approximate Monte Carlo algorithms to estimate result probabilities. We scale our algorithm to large scale junction trees using the previously proposed INDSEP data structure. Further, we develop novel optimization techniques to process a batch of lineages by sharing computation across formulae that involve common subexpressions, and adapt our algorithms to exploit any independence relationships that exist in the data. Our experimental study illustrates the benefits of using our system for processing lineage formulae over correlated probabilistic databases.

Monte Carlo Processing of Probabilistic Satisfiability Queries in MCDB

Luis Perez (Rice University), Subi Arumugam (University of Florida), Christopher Jermaine (Rice University)

MCDB is a prototype database system for managing stochastic models for uncertain data. Because MCDB allows a user to specify arbitrary models for the uncertain data that are hard to analyze directly, MCDB utilizes a purely sampling-based approach to query processing. To evaluate a query, MCDB uses the underlying stochastic models to pseudorandomly generate many possible worlds, and then it runs the query on each of them in order to measure the effect of the uncertainty embodied in the stochastic model upon the query result. In this paper, we study the problem of how to use MCDB to answer statistical queries that search for database objects which satisfy some filter condition with greater (or less than) a user-specified probability. For example: ``Which packages will arrive late with > 5% probability?" ``Which regions will see more than a 2% decline in sales with > 50% probability?" ``What items will be out of stock by Friday with > 20% probability?" We consider both the systems aspects and the statistical aspects of the problem.

K-Nearest Neighbor Search for Fuzzy Objects

Kai Zheng, Pui Cheong Fung, Xiaofang Zhou (University of Queensland)

The k nearest neighbor (kNN) problem has long been an attractive and flourishing topic due to its broad range of applications. In this work we try to study this problem in the context that both the target data and query are fuzzy, i.e. they have indetermined boundaries. Fuzzy objects play an important role in biomedical image databases, GIS systems and so on. However, existing work mainly focus on mathematical modelling such as defining basic fuzzy types and operations, leaving more advanced quries such as kNN query untouched. In this paper, we firstly define a distance function for measuring the spatial closeness of two fuzzy objects, while at the same time satisfying the confidence level that users wish to be. Upon this new measurement, we then define two kinds of queries, \emph{Ad-hoc kNN query} (AKNN) and \emph{Range kNN query} (RKNN), which benefit users in looking for k closest objects at desired probability levels. For efficient evaluation of AKNN query, we utilize R-tree as index to reduce the number of unnecessary object probe by invoking a cheap check of lower and upper bounds before calculating the actual distance. Tighter bounds based on the observation of ``shrinking" property of objects are also developed to further improve the pruning power. To answer RKNN query efficiently, a series of lemmas are proposed, inspired by which our algorithm can filter out most disqualified objects, identify results as early as possible and avoid exhaustively scan of entire range simutaneously. Finally the effiency of our algorithms are verified by extensive experiments on synthetic and real datasets, and the algorithms achieve satisfactory performance.

An Optimal Labeling Scheme for Workflow Provenance Using Skeleton Labels

Zhuowei Bao, Susan Davidson, Sanjeev Khanna, Sudeepa Roy (University of Pennsylvania)

We develop a compact and efficient reachability labeling scheme for answering provenance queries on workflow runs that conform to a given specification. Even though a workflow run can be structurally more complex and can be arbitrarily larger than the specification due to fork (parallel) and loop executions, we show that a compact reachability labeling for a run can be efficiently computed using the fact that it originates from a fixed specification. Our labeling scheme is optimal in the sense that it uses labels of logarithmic length, runs in linear time, and answers any reachability query in constant time. Our approach is based on using the reachability labeling for the specification as an effective "skeleton" for designing the reachability labeling for workflow runs. We also demonstrate empirically the effectiveness of our skeleton-based labeling approach.

SIGMOD Research Session 16: Data Security & Privacy

Chair: Chris Clifton (Purdue University)

Thursday, 10:30 – 12:00

SecureBlox: Customizable Secure Distributed Data Processing

William Marczak (UC Berkeley), Shan Shan Huang, Martin Bravenboer (LogicBlox, Inc), Micah Sherr, Boon Thau Loo (University of Pennsylvania), Molham Aref (LogicBlox, Inc)

In this paper, we present SecureBlox, a declarative system that unifies a distributed query processor with a security policy framework. In SecureBlox, programmers compose existing mechanisms to compactly specify and reconfigure security policies. Our implementation of SecureBlox is a series of extensions to LogicBlox, an emerging commercial Datalog-based platform for enterprise software systems, with enhancements to enable distribution, integrity constraints and static meta-programmability. SecureBlox achieves meta programmability via BloxGenericsL: a language extension that provides compile-time code generation based on the security requirements and trust policies of the deployed environment. We present and evaluated detailed use-cases where SecureBlox enables applications such as an authenticated declarative routing protocol with encrypted advertisements and an authenticated and encrypted parallel hash join operation. Our results demonstrate SecureBlox's ability to specify and implement a wide range of different security constructs for distributed systems, and enable tradeoffs between performance and security.

Differentially Private Aggregation of Distributed Time-Series with Transformation and Encryption

Vibhor Rastogi (University of Washington), Suman Nath (Microsoft Research)

We propose the first differentially private aggregation algorithm for distributed time-series data that offers good practical utility without any trusted server. This addresses two important challenges in participatory data-mining applications where (i) individual users wish to publish temporally correlated time-series data (such as location traces, web history, personal health data), and (ii) an untrusted third-party aggregator wishes to run aggregate queries on the data. To ensure differential privacy for time-series data despite the presence of temporal correlation, we propose the Fourier Perturbation Algorithm (FPA). Standard differential privacy techniques perform poorly for time-series data. To answer n queries, such techniques can result in a noise of Theta(n) to each query answer, making the answers practically useless if n is large. Our FPA perturbs the Discrete Fourier Transform of the query answers. For answering n queries, FPA improves the expected error from Theta(n) to roughly Theta(k) where k is the number of Fourier coefficients that can (approximately) reconstruct all the n query answers. Our experiments show that k << n for many real-life data-sets resulting in a huge error-improvement for FPA. To deal with the absence of a trusted central server, we propose the Distributed Laplace Perturbation algorithm~(DLPA) to add noise in a distributed way in order to guarantee differential privacy. To the best of our knowledge, DLPA is the first distributed differentially private algorithm that can scale with a large number of users: DLPA outperforms the only other distributed solution for differential privacy proposed so far, by reducing the computational load per user from O(U) to O(1) where U is the number of users.

Non-homogeneous Generalization in Privacy Preserving Data Publishing

Wai Kit Wong, Nikos Mamoulis, David Cheung (University of Hong Kong)

Previous research on privacy-preserving data publishing, based on the k-anonymity model, has followed the simplistic approach of homogeneously giving the same generalized value in all quasi-identifiers within a partition. We observe that the anonymization error can be reduced if we follow a non-homogeneous generalization approach for groups of size larger than k. Such an approach would allow tuples within a partition to take different generalized quasi-identifier values. Anonymization following this model is not trivial, as its direct application can easily violate k-anonymity. In addition, non-homogeneous generalization allows for additional types of attack, which should be considered in the process. We provide a methodology for verifying whether a non-homogeneous generalization violates k-anonymity. Then, we propose a technique that generates a non-homogeneous generalization for a partition and show that its result satisfies k-anonymity, however

by straightforwardly applying it, privacy can be compromised if the attacker knows the anonymization algorithm. Based on this, we propose a randomization method that prevents this type of attack and show that k-anonymity is not compromised by it. Non-homogeneous generalization can be used on top of any existing partitioning approach to improve its utility. In addition, we show that a new partitioning technique tailored for non-homogeneous generalization can further improve quality. A thorough experimental evaluation demonstrates that our methodology greatly improves the utility of anonymized data in practice.

Preserving Privacy and Fairness in Peer-to-Peer Data Integration

Hazem Elmeleegy, Mourad Ouzzani, Ahmed Elmagarmid, Ahmad Abusalah (Purdue University)

Peer-to-peer data integration - a.k.a. Peer Data Management Systems (PDMSs) - promises to extend the classical data integration approach to the Internet scale. Unfortunately, some challenges remain before realizing this promise. One of the biggest challenges is preserving the privacy of the exchanged data while passing through several intermediate peers. Another challenge is protecting the mappings used for data translation. Protecting the privacy without being unfair to any of the peers is yet a third challenge. This paper presents a novel query answering protocol in PDMSs to address these challenges. The protocol employs a technique based on noise selection and insertion to protect the query results, and a commutative encryption-based technique to protect the mappings and ensure fairness among peers. An extensive security analysis of the protocol shows that it is resilient to seven possible types of attacks. We implemented the protocol within an established PDMS: the Hyperion system. We conducted an experimental study using real data from the healthcare domain. The results show that our protocol manages to achieve its privacy and fairness goals, while maintaining query processing time at the interactive level.

SIGMOD Research Session 17: Web Data Integration

Chair: Fatma Ozcan (IBM Almaden)

Thursday, 10:30 – 12:00

Structured Annotations of Web Queries

Nikos Sarkas (University of Toronto), Stelios Paparizos, Panayiotis Tsaparas (Microsoft Research)

Queries asked on web search engines often target information residing in structured data sources. For instance, structured data can be used to provide answers to queries about commercial products (e.g. `canon eos 350d'), movie show-times (e.g.`indiana jones 4 near boston') or airline schedules (`flights from boston to new york'). However, surfacing relevant results from such data poses important challenges: (i) answers should be provided within fraction of a second, (ii) large heterogeneous collections of data exists, (iii) users targeting structured data express queries in unstructured free-form text without knowledge of schema and (iv) keyword queries targeting structured data used the same language as other web queries. In this paper we address the problem of discovering latent structured semantics in web queries to produce \emph{Structured Annotations} for them. We consider an annotation as a mapping from a query to a table of structured data and attributes of this table. Given a collection of structured tables, we present a fast and scalable tagging mechanism for obtaining all possible annotations of a query over these tables. However, we observe that for a given query only few are sensible for the user needs. We thus propose a principled probabilistic scoring mechanism using a generative model which sorts the produced annotations based on the likelihood of mapping the query to a given table and its attributes. Furthermore, we extend the model to create a dynamic threshold, filtering out misinterpreted query annotations. Our techniques are completely unsupervised, obviating the need for costly manual labeling effort. We evaluated our techniques using real world web queries and structured data present very promising results.

On Active Learning of Record Matching Packages

Arvind Arasu (Microsoft Research), Michaela Goetz (Cornell University), Raghav Kaushik (Microsoft Research)

We consider the problem of learning a record matching package (classifier) in an active learning

setting. In active learning, the learning algorithm picks the set of examples to be labeled, unlike more traditional passive learning where a user selects the labeled examples. Active learning is important for record matching since manually identifying a suitable set of labeled examples is difficult. Previous algorithms that use active learning for record matching have serious limitations: The packages that they learn lack quality guarantees and the algorithms do not scale to large input sizes. We present new algorithms for this problem that overcome these limitations. Our algorithms are fundamentally different from traditional active learning approaches, and are designed ground up to exploit problem characteristics specific to record matching. Our algorithms can also be adapted to work in the passive learning setting. We include a detailed experimental evaluation on real-world data demonstrating the effectiveness of our algorithms.

I4E: Interactive Investigation of Iterative Information Extraction

Anish Das Sarma (Yahoo Research), Alpa Jain (Yahoo), Divesh Srivastava (AT&T Labs-Research)

Information extraction systems are increasingly being used to mine structured information from unstructured text documents. A commonly used unsupervised technique is to build {\em iterative information extraction (IIE)} systems that learn task-specific rules, called {\em patterns}, to generate the desired tuples. Oftentimes, output from an information extraction system may contain unexpected results which may be due to an incorrect pattern, incorrect tuple, or both. In such scenarios, users and developers of the extraction system could greatly benefit from an investigation tool that can quickly help them reason about and repair the output. In this paper, we develop an approach for interactive post-extraction investigation for IIE systems. We formalize three important phases of this investigation, namely, {\em explain} the IIE result, {\em diagnose} the influential and problematic components, and {\em repair} the output from an information extraction system. We show how to characterize the execution of an IIE system and build a suite of algorithms to answer questions pertaining to each of these phases. We experimentally evaluate our proposed approach over several domains over a Web corpus of about 500 million documents. We show that our approach effectively enables post-extraction investigation, while maximizing the gain from user and developer interaction.

ONDUX: On-Demand Unsupervised Learning for Information Extraction

Eli Vilarinho (Federal University of Amazonas), Altigran Silva (UFAM), Marcos Goncalves (UFMG), Edleno de Moura (Federal University of Amazonas)

Information extraction by text segmentation (IETS) applies to cases in which data values of interest are organized in implicit semi-structured records available in textual sources (e.g. postal addresses, bibliographic information, ads). It is an important practical problem that has been frequently addressed in the recent literature. In this paper we introduce ONDUX (On Demand Unsupervised Information Extraction), a new unsupervised probabilistic approach for IETS. As other unsupervised IETS approaches, ONDUX relies on information available on pre-existing data to associate segments in the input string with attributes of a given domain. Unlike other approaches, we rely on very effective matching strategies instead of explicit learning strategies. The effectiveness of this matching strategy is also exploited to disambiguate the extraction of certain attributes through a reinforcement step that explores sequencing and positioning of attribute values directly learned {\em on-demand} from test data, with no previous human-driven training, a feature unique to ONDUX. This assigns to ONDUX a high degree of flexibility and results in superior effectiveness, as demonstrated by the experimental evaluation we report with textual sources from different domains, in which ONDUX is compared with a state-of-art IETS approach.

SIGMOD Research Session 18: Web Data Management

Chair: Jun Tatemura (NEC Research Laboratories)

Thursday, 13:30 – 15:00

Optimizing Content Freshness of Relations Extracted From the Web Using Keyword Search

Mohan Yang (Shanghai Jiao Tong University), Haixun Wang (Microsoft Research, Asia), Lipyeow Lim, Min Wang (HP Labs)

An increasing number of applications operate on data obtained from the Web. These applications typically maintain local copies of the Web data to avoid network latency in data accesses. As the data on the Web evolves, it is critical that the local copy be kept up-to-date. Data freshness is one of the most important data quality issues in information systems, and has been extensively studied for various applications including Web crawling. Web crawling is focused on obtaining as many raw web pages as possible. Our applications, on the other hand, are interested in specific content from specific data sources. Knowing the content or the semantics of the data enables us to differentiate data items based on their importance and volatility, which are key factors that impact the design of the data synchronization strategy. In this work, we formulate the concept of content freshness, and present a novel approach that maintains content freshness with least amount of Web communication. Specifically, we assume data is accessible through a general keyword search interface, and we form keyword queries based on their selectivity, as well their contribution to content freshness of the local copy. Experiments show the effectiveness of our approach compared with several naive methods for keeping data fresh.

Feeding Frenzy: Selectively Materializing Users' Event Feeds

Adam Silberstein (Yahoo! Research), Jeffrey Terrace (Princeton University), Brian Cooper, Raghu Ramakrishnan (Yahoo! Research)

Near real-time event feeds are becoming a key feature of many popular web applications. Examples include user generated events on Twitter and Facebook, and news stories on particular topics on iGoogle and My Yahoo. How can we efficiently construct a web page showing the latest events from a user's feed? Constructing such a feed must be fast so the page loads quickly. However, the wide fanout of events (some sources have many followers) and high skew (fanout and rates vary widely) make it difficult to scale such applications. We argue that the best performance results from selectively materializing each user's feed: events from high-rate producers are retrieved at query time, while events from lower-rate producers are materialized in advance. Furthermore, a formal analysis of the problem shows that the proper strategy depends on the ratio between a given producer's event rate and a given consumer's page view rate. Thus, some producers may be materialized for some consumers and not others; and portions of a consumer's feed may be materialized while other portions are not. Our experimental results, using a real web database infrastructure, shows that this hybrid strategy results in the lowest system load (and hence improves scalability) under a variety of workloads.

Constructing and Exploring Composite Items

Senjuti Basu Roy (University of Texas at Arlington), Sihem Amer-Yahia (Yahoo! Research), Ashish Chawla (Yahoo! Inc), Gautam Das (University of Texas at Arlington), Cong Yu (Yahoo! Research)

Nowadays, online shopping has become a daily activity. Web users purchase a variety of items ranging from books to electronics. The large supply of online products calls for sophisticated techniques to help users explore available items. We propose to build {\em composite items} which associate a {\em central item} with {\em a set of packages}, formed by {\em satellite items}, and help users explore them. For example, a user shopping for an {\em iPhone} (i.e., the central item) with a price budget can be presented with both the iPhone and a package of other items that matches well with the iPhone (e.g., \{{\em Belkin case, Bose sounddock, Kroo USB cable}\}) as a composite item, whose total price is within the user's budget. We define and study the problem of effective construction and exploration of large sets of packages associated with a central item, and design and implement efficient algorithms for solving the problem in two stages: {\em summarization}, a technique which picks k representative packages for each central item; and {\em visual effect optimization}, which helps the user find diverse composite items quickly by minimizing overlap between packages presented to the user in a ranked order. We conduct an extensive set of experiments on real life data sets to demonstrate the efficiency and effectiveness of our algorithms.

Unbiased Estimation of Size and other Aggregates over Hidden Web Databases

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

Many websites provide restrictive form-like interfaces which allow users to execute search queries on the underlying hidden databases. In this paper, we consider the problem of estimating the size of a hidden database through its web interface. We propose novel techniques which use a small number of queries to produce unbiased estimates with small variance. These techniques can also be used for approximate query processing over hidden databases. We present theoretical analysis and extensive experiments to illustrate the effectiveness of our approach.

SIGMOD Research Session 19: Graph Mining

Chair: Chen Li (UC Irvine)

Thursday, 13:30 – 15:00

Towards Proximity Pattern Mining in Large Graphs

Arijit Khan, Xifeng Yan, Kun-Lung Wu (IBM Watson Research Center)

Mining graph patterns in large networks is critical to a variety of applications such as malware detections and biological module discovery. However, frequent subgraphs are often ineffective to capture association existing in these applications, due to the complexity of isomorphism testing and the inelastic pattern definition. In this paper, we introduce proximity pattern which is a signifcant departure from the traditional concept of frequent subgraphs. Defined as a set of labels that co-occur in neighborhoods, proximity pattern blurs the boundary between itemset and structure. It relaxes the rigid structure constraint of frequent subgraphs, while introducing structure association to frequent itemsets. Therefore, it could benefit from both: efficient mining in itemsets and connectivity from graphs. We propose an information propagation model, called Normalized Probabilistic Association NmPA) to transform the mining problem of proximity patterns to a weighted association mining problem. NmPA is evaluated on real-life social and intrusion networks. Empirical results show that it not only finds interesting patterns that are ignored by the existing approaches, but also achieves high performance to find proximity patterns in large scale graphs.

GAIA: Graph Classification Using Evolutionary Computation

Ning Jin, Calvin Young, Wei Wang (University of North Carolina at Chapel Hill)

Discriminative subgraphs are widely used to define the feature space for graph classification in large graph databases. Several scalable approaches have been proposed to mine discriminative subgraphs. However, their intensive computation need prevents them from mining large databases. We propose an efficient method GAIA for mining discriminative subgraphs for graph classification in large databases. Our method employs a novel subgraph encoding approach to allow arbitrary subgraph pattern exploration order and explores subgraph pattern space in a process resembling biological evolution. In this manner, GAIA is able to find discriminative subgraph patterns much faster than other algorithms. Additionally, we take advantage of parallel computing to further improve the quality of resulting patterns. In the end, we employ sequential coverage to generate association rules as graph classifiers using patterns mined by GAIA. Extensive experiments have been performed to analyze the performance of GAIA and to compare it with two other state-of-the-art approaches. GAIA outperforms the other approaches in terms of both the classification accuracy and runtime efficiency.

Finding Maximum Degrees in Hidden Bipartite Graphs

Yufei Tao, Sheng Cheng (Chinese University of Hong Kong), Jianzhong Li (Harbin Institute of Technology)

A {\em hidden graph} is a graph whose edges are not explicitly given. Detecting the presence of an edge requires an expensive {\em edge-probing} query. We consider the {\em k most connected vertex} problem on hidden bipartite graphs. Specifically, given a bipartite graph G with independent vertex sets B and W, the goal is to find the k vertices in B with the largest

degrees by performing the minimum number of queries. This problem is encountered in many applications such as query optimization in a DBMS, data mining, marketing research, and so on. If B and W have n and m vertices respectively, the number of queries needed to solve the problem is nm in the worst case. This, however, is a pessimistic estimate on how many queries are necessary on practical data. In fact, on some inputs, the problem can be settled with at most $km + n$ edges, which is significantly lower than nm for $k \ll n$. It is thus interesting to design an {\em adaptive} algorithm that is guaranteed to achieve the best possible performance on every input G. We give such an algorithm, and prove that it is {\em instance optimal} among a broad class of solutions. This means that, for {\em any} G, our algorithm can perform more queries than the optimal solution (which is currently unknown) by only a constant factor, which can be shown to be at most 2. Furthermore, when $k = O(1)$ (e.g., 10), the proposed algorithm can be worse than the optimal solution by only a tiny factor of $1 + O(1/n)$. Extensive experiments demonstrate that, in practice, the number of queries our technique performs is far less than nm, and agrees with our theoretical findings very well.

Connected Substructure Similarity Search

Haichuan Shang, Xuemin Lin, Wei Wang (University of New South Wales), Jeffrey Xu Yu (Chinese University of Hong Kong), Ying Zhang (University of New South Wales)

Substructure similarity search is to retrieve graphs that approximately contain a given query graph. It has many applications, e.g., detecting similar functions among chemical compounds. The problem is challenging as even testing subgraph containment between two graphs is NP-complete. Hence, existing techniques adopt the filtering-and-verification framework with the focus on developing effective and efficient techniques to remove non-promising graphs. Nevertheless, existing filtering techniques may be still unable to effectively remove many ``low" quality candidates. To resolve this, in this paper we propose a novel indexing technique, GrafD-Index, to index graphs according to their ``distances" to features. We characterize a tight condition under which the distance-based triangular inequality holds. We then develop lower and upper bounding techniques that exploit the GrafD-Index to (1) prune non-promising graphs and (2) include graphs whose similarity is guaranteed to exceed the given similarity threshold. Considering that the verification phase is not well studied and plays the dominant role in the whole process, we devise efficient algorithms to verify candidates. A comprehensive experiment using real datasets demonstrates that our proposed method techniques significantly outperforms existing methods.

SIGMOD Research Session 20: Indexing & Storage Management

Chair: Daniel Abadi (Yale University)

Thursday, 13:30 – 15:00

B^ed-Tree: An All-Purpose Tree Index for String Similarity Search on Edit Distance

Zhenjie Zhang, Beng Chin Ooi (National University of Singapore), Marios Hadjieleftheriou, Divesh Srivastava (AT&T Labs – Research)

Strings are ubiquitous in computer systems, attracting extensive research efforts from computer scientists in different areas. One of the most important problems on string processing is how to efficiently conduct similarity search on strings with respect to some distance measure, with applications on information retrieval, data cleaning and biological sequence analysis. While a couple of different distance functions on strings have been proposed, edit distance is the most natural option which reflects how many basic operations are necessary to change one string to another. Surprisingly, the scalability remains unsatisfactory with the inverted list structure to support string search on edit distance, especially in scenarios with strict memory constraint or frequent updates. In this paper, we present a general and simple framework for string processing on the standard B+ tree structure, with carefully designed transformation from string domain to integer space. We show that this all-purpose index scheme supports different queries, including range query, top-k query and join query, only depending on the properties of the engaged mapping function. Three mapping functions are thus discussed, capturing useful information in prefix, gram statistics and mismatch locations of the indexed strings respectively. Extensive experiments on real data sets imply that our proposal

greatly improves the scalability and achieves almost the same querying performance as the state-of-the-art methods in the literature do.

On Indexing Error-Tolerant Set Containment

Raghav Kaushik (Microsoft Research), Parag Agrawal (Stanford University), Arvind Arasu (Microsoft Research)

Prior work has identified set based comparisons as a useful primitive for supporting a wide variety of similarity functions in record matching. Accordingly, various techniques have been proposed to improve the performance of set similarity lookups. However, this body of techniques focuses almost exclusively on {\em symmetric} notions of set similarity. In this paper, we study the indexing problem for the asymmetric {\em jaccard containment} similarity function that is an error-tolerant variation of set containment. We enhance this similarity function to also account for string transformations that reflect synonyms such as ``Mike'' and ``Michael'' referring to the same first name. To our knowledge, this is the first paper that studies this indexing problem.

Workload-Aware Storage Layout for Database Systems

Oguzhan Ozmen, Kenneth Salem (University of Waterloo), Jiri Schindler, Steve Daniel (NetApp, Inc)

The performance of a database system depends strongly on the layout of database objects, such as indexes or tables, onto the underlying storage devices. A good layout will both balance the I/O workload generated by the database system and avoid the performance-degrading interference that can occur when concurrently accessed objects are stored on the same volume. In current practice, layout is typically guided by heuristics and rules of thumb, such as separating indexes and tables or striping all objects across all of the available storage devices. However, these guidelines may not always give good results. In this paper, we address the problem of generating an optimized layout of a given set of database objects. Our layout optimizer goes beyond generic guidelines by making use of a description of the database system's I/O activity. We formulate the layout problem as a non-linear programming (NLP) problem and use the I/O description as input to an NLP solver. Our layout optimization technique, which is incorporated into a database layout advisor, identifies a layout that both balances load and avoids interference. We evaluate experimentally the efficacy of our approach and demonstrate that it can quickly identify non-trivial optimized layouts.

Querying Data Provenance

Grigoris Karvounarakis, Zachary Ives, Val Tannen (University of Pennsylvania)

Many advanced data management operations (e.g., incremental maintenance, trust assessment, debugging schema mappings, keyword search over databases, or query answering in probabilistic databases), involve computations that look at how a tuple was produced, e.g., to determine its score or existence. This requires answers to queries such as, ``Is this data derivable from trusted tuples?''; ``What tuples are derived from this relation?''; or ``What score should this answer receive, given initial scores of the base tuples?''. Such questions can be answered by consulting the provenance of query results. In recent years there has been significant progress on formal models for provenance. However, the issues of provenance storage, maintenance, and queries have not yet been addressed in an application-independent way. In this paper, we adopt the most general formalism for tuple-based provenance, semiring provenance. We develop a query language for provenance, which can express all of the aforementioned types of queries, as well as many more; we propose storage, processing and indexing schemes for data provenance in support of these queries; and we experimentally validate the feasibility of provenance querying and the benefits of our indexing techniques across a variety of application classes and queries.

SIGMOD INDUSTRIAL PAPER ABSTRACTS

SIGMOD Industrial Session 1: New Platforms

Chair: Divy Agrawal (UC Santa Barbara)
Tuesday 10:30 – 12:00

Experiences Evolving a New Analytical Platform: What Works and What's Missing
Jeff Hammerbacher (Cloudera)

At Cloudera, we augment existing analytical platforms with some new tools for data management and analysis. In this talk, we'll share some experiences of what has worked across industries and workloads, and what new software components might help complete a new analytical platform.

Overview of SciDB: Large Scale Array Storage, Processing and Analysis
Paul Brown (SciDB)

SciDB is a new open-source data management system intended primarily for use in application domains that involve very large (Petabyte) scale array data; for example, scientific applications such as astronomy, remote sensing and climate modeling, bio-science information management, risk management systems in financial applications, and the analysis of web log data. In this talk we will describe our set of motivating examples and use them to explain the features of SciDB. We then briefly give an overview of the project 'in flight', explaining our novel storage manager, array data model, query language, and extensibility frameworks.

Integrating Hadoop and parallel DBMS
Yu Xu, Pekka Kostamaa, Like Gao (Teradata)

Teradata's parallel DBMS has been successfully deployed in large data warehouses over the last two decades for large scale business analysis in various industries over data sets ranging from a few terabytes to multiple petabytes. However, due to the explosive data volume increase in recent years at some customer sites, some data such as web logs and sensor data are not managed by Teradata, partially because it is very expensive to load those extreme large volumes of data to a DBMS such as Teradata, especially when those data are not frequently used to support important business decisions. Recently the MapReduce programming paradigm, started by Google and made popular by the open source Hadoop implementation with major support from Yahoo, is gaining rapid momentum in both academia and industry as another way of performing large scale data analysis. By now most BI analysts and data warehouse practitioners agree that both parallel DBMS and MapReduce paradigms have advantages and disadvantages for various business applications and thus both paradigms are going to coexist for a long time. In fact, a large number of Teradata customers, especially those in the e-business and telecom industries have seen increasing needs to perform BI over both data stored in Hadoop and data in Teradata EDW (Enterprise Data Warehouse). Clearly transferring data between Hadoop and Teradata EDW efficiently is the important first step for integrated BI over Hadoop and Teradata EDW. A straightforward approach without the need of any new development from either the Hadoop or Teradata side is to use Hadoop and Teradata's current load and export utilities: HDFS files can be copied to regular files which can be loaded to Teradata, and tables from Teradata can be exported to files which can be loaded to HDFS. However, one common thing between Hadoop and Teradata EDW is that data in both systems are partitioned across multiple nodes for parallel computing, which creates optimization opportunities not possible for DBMSes running on a single node. We describe our three major efforts towards tight and efficient integration of Hadoop and Teradata EDW.

A Comparison of Join Algorithms for Log Processing in MapReduce
Spyros Blanas, Jignesh Patel (University of Wisconsin), Vuk Ercegovac, Jun Rao, Eugene Shekita, Yuanyuan Tian (IBM Almaden Research Center)

The MapReduce framework is increasingly being used to analyze large volumes of data. One important type of data analysis done with MapReduce is log processing, in which a click-stream or an event log is filtered, aggregated, or mined for patterns. As part of this analysis, the log often needs to be joined with reference data such as information about users. Although there have been many studies examining join algorithms in parallel and distributed DBMSs, the MapReduce framework is cumbersome for joins. MapReduce programmers often use simple but ine±cient algorithms to perform joins. In this paper, we describe crucial implementation details of a number of well-known join strategies in MapReduce, and present a comprehensive experimental comparison of these join techniques on a 100- node Hadoop cluster. Our results provide insights that are unique to the MapReduce platform and offer guidance on when to use a particular join algorithm on this new platform.

SIGMOD Industrial Session 2: Advanced Analytics

Chair: Berthold Reinwald (IBM Almaden Research Center)
Tuesday 13:30 – 15:00

Ricardo: Integrating R and Hadoop

Yannis Sismanis (IBM Almaden Research Center), Sudipto Das (UC Santa Barbara), Rainer Gemulla, Peter Haas, Kevin Beyer, John McPherson (IBM Almaden Research Center)

Many of today's enterprises are collecting data at the most detailed level possible, thereby creating data repositories from the terabyte to petabyte range. The ability to deeply analyze this data is becoming increasingly essential for be- ing competitive in the market. Yet, this interplay of large data and deep analysis poses a significant challenge to ex- isting statistical software and data management systems. On the one hand, statistical software provides rich func- tionality for data analysis and modeling, but is designed to handle only limited amounts of data; e.g., popular packages like R and SPSS operate entirely in main memory. On the other hand, data management systems|such as analyt- ical databases or MapReduce-based systems|have proven to scale to petabytes of data, but they do not provide sufficient statistical functionality. We report our experiences in combining both systems in order to build RICARDO, a scal- able platform for deep analytics in the cloud. RICARDO is named after David Ricardo, a famous economist, and rests on a decomposition of statistical algorithms into parts ex- ecuted by a statistical software package and parts handled by a cloud-enabled data management system. This contrasts previous approaches|which try to get along with only one of the systems|in that suboptimal transfer of functionality across system boundaries is avoided. Our results are vali- dated by a case study in which we scale one of the winning algorithms of the recent Net ix competition to handle bil- lions of movie ratings and about a terabyte of data using a cluster of commodity machines.

PYMK: Friend Recommendation at MySpace

Michael Moricz, Yerbolat Dosbayev, Mikhail Berlyant (MySpace.com)

In recent years Social Networking has enjoyed a significant increase in popularity. The main reason behind this surge in popularity is the social experience associated with connecting content to people and also connecting people with other people. Knowing, seeing, hearing what our friends and like-minded people feel or listen to or upload is an unparalleled experience. Similar to real life, finding good friends is not easy without the help of good recommendations. In this Extended Abstract we present the MySpace friend recommendation algorithm named People You May Know. We will also comment on both the quality and the effectiveness of the algorithms.

Forecasting High-Dimensional Data

Deepak Agarwal (Yahoo! Research), Datong Chen, Long-ji Lin (Yahoo! Labs), Jayavel Shanmugasundaram, Erik Vee (Yahoo! Research)

We propose a method for forecasting high-dimensional data (hundreds of attributes, trillions of attribute combinations) for a duration of several months. Our motivating application is guaranteed display advertising, a multi-billion dollar industry, whereby advertisers can buy targeted (high-

dimensional) user visits from publishers many months or even years in advance. Forecasting high-dimensional data is challenging because of the many possible attribute combinations that need to be forecast. To address this issue, we propose a method whereby only a {\em sub-set} of attribute combinations are explicitly forecast and stored, while the other combinations are dynamically forecast on-the-fly using high-dimensional attribute correlation models. We evaluate various attribute correlation models, from simple models that assume the independence of attributes to more sophisticated sample-based models that fully capture the correlations in a high-dimensional space. Our evaluation using real-world display advertising data sets shows that fully capturing high-dimensional correlations leads to significant forecast accuracy gains. The proposed method is in production use in Yahoo!'s guaranteed delivery display advertising system.

Datawarehousing and Analytics Infrastructure at Facebook

Ashish Thusoo, Dhruba Borthakur (Facebook)

Scalable analysis on large data sets has been core to the functions of a number of teams at Facebook - both engineering and non-engineering. Apart from ad hoc analysis and business intelligence applications used by analysts across the company, a number of Facebook products are also based on analytics. These products range from simple reporting applications like Insights for the Facebook Ad Network, to more advanced kinds such as Facebook's Lexicon product. As a result a flexible infrastructure that caters to the needs of these diverse applications and users and that also scales up in a cost effective manner with the ever increasing amounts of data being generated on Facebook, is critical. Hive and Hadoop are the technologies that we have used to address these requirements at Facebook. In this talk we will present the general architecture of log collection and datawarehousing at facebook and also introduce core building blocks that we have used to implement a peta byte scale datawarehousing solution.

SIGMOD Industrial Session 3: Advances in DBMSs

Chair: Sunil Prabhakar (Purdue University)

Wednesday 8:30 – 10:00

Extreme Scale with Full SQL Language Support in Microsoft SQL Azure

Nigele Ellis, Gopal Kakivaya , Dave Campbell (Microsoft)

Cloud SQL Server is an Internet scale relational database service which is currently used by Microsoft delivered services and also offered directly as a fully relational database service known as “SQL Azure”. One of the principle design objectives in Cloud SQL Server was to provide true SQL support with full ACID transactions within controlled scale “consistency domains” and provide a relaxed degree of consistency across consistency domains that would be viable to clusters of 1,000’s of nodes. In this paper, we describe the implementation of Cloud SQL Server with an emphasis on this core design principle.

Pay-As-You-Go - an adaptive approach to provide full context aware text search over document content

Zhen Hua Liu, Thomas Baby, Sukhendu Chakraborty, Junyan Ding, Anguel Novoselsky, Vikas Arora (Oracle)

RDBMS provides best performance for querying structured data that starts out with a well-defined schema. However, such a ‘schema first, data later’ approach does not work for unstructured data or data without much structure. Therefore, RDBMS typically stores such data without any schema in LOB columns (for example, Character Large Object (CLOB) or Binary Large Object (BLOB) columns) and provides Information-Retrieval (IR) style, keyword-based search capability over these LOB columns. Lately, XML as a native datatype (XMLType) in RDBMS has been introduced via the SQL/XML standard. Semi-structured data with or without any schema can be stored into such XMLType columns, and Xquery provides query capability over them. In particular, XQuery full text specification provides the capability of searching keywords within document context. Such full context-aware text search capability is more powerful than pure keyword search, since the user can now provide fine-grained context in which the keywords should occur. However, XML with XQuery

full text searching requires that the user first convert her text data into XML and store them into XMLType column. Such massive physical data migration with possible loss of document fidelity and its potential impact on existing production environments are often expensive enough that users are reluctant to adopt the XML/XQuery approach. In this paper, we propose a pay-as-you-go architecture to provide XML text view over LOB columns, so that user can take advantage of context-aware full-text search capability adaptively. This adaptive architecture includes a novel XML text index that can be created over the LOB column where the content is stored. The XML text index supports an XML text view over LOB data on top of which XQuery full-text search capability is feasible. Such an adaptive index/view approach provides least intrusion over existing data, as it requires no physical data migration. We describe the design and challenge of building such an adaptive XML text index. Furthermore, we advocate that the pay-as-you-go approach provides the integration bridge between the structured relational world and text oriented document world and fulfils the primary motivation of XML in the database.

Sedna: Native XML Database Management System (Internals Overview)

Dmitry Lizorkin (Institute for System Programming of the Russian Academy of Sciences)

We present a native XML database management system, Sedna, which is implemented from scratch as a full-featured database management system for storing large amounts of XML data. We believe that the key contribution of this system is an improved schema-based clustering storage strategy efficient for both XML querying and updating, and powered by a novel memory management technique. We position our approach with respect to state-of-the-art methods, describe the main techniques and demonstrate a wide range of queries and updates issued on massive data sets to illustrate the pros and cons of our approach.

Optimizing Tuple-store Query Execution

Scott Meyer, Jutta Degener, John Giannandrea, Barak Michener (Metaweb Technologies Inc.)

Relational databases require that a database schema exist prior to data entry and require manual optimization (via CREATE INDEX) and tuning for best performance. We describe the query optimization techniques used by graphd, a schema-last, automatically indexed tuple-store which supports freebase.com, a world-writable database. Graphd is a log-structured store with a query optimizer based on a functional operator tree over the domain of sorted integer sets, sets which arise naturally as tuples are appended to the store. We demonstrate that a set-based optimizer can deliver performance that is roughly comparable to traditional RDBMS query optimization techniques applied to a fixed schema.

SIGMOD Industrial Session 4: Information Integration

Chair: Chen Li (UC Irvince)

Wednesday 10:30 – 12:00

OpenII: An Open Source Information Integration Toolkit

Len Seligman, Peter Mork (MITRE), Alon Halevy (Google), Ken Smith (MITRE), Michael Carey (UC Irvine), Kuang Chen (UC Berkeley), Chris Wolf (MITRE), Jayant Madhavan (Google), Akshay Kannan (UC Berkeley)

OpenII (openintegration.org) is a collaborative effort to create a suite of open-source tools for information integration (II). The project is leveraging the latest developments in II research to create a platform on which integration tools can be built and further research conducted. In addition to a scalable, extensible platform, OpenII includes industrial-strength components developed by MITRE, Google, UC-Irvine, and UC-Berkeley that interoperate through a common repository in order to solve II problems. Components of the toolkit have been successfully applied to several large-scale US government II challenges.

Google Fusion Tables: Data Management, Integration and Collaboration in the Cloud

Jonathan Goldberg-Kidon, Hector Gonzalez, Alon Halevy, Christian Jensen, Anno Langen, Jayant Madhavan, Rebecca Shapely (Google Inc.)

Google Fusion Tables is a cloud-based service for data management and integration. The service was launched in June, 2009, and has since received a considerable amount of use (see tables.googlelabs.com). Though we have witnessed a wide range of applications of Fusion Tables, the audience for which the service was designed is organizations that are struggling with making their data available internally and externally, and communities of users that need to collaborate on data management across multiple enterprises. Fusion Tables enables users to upload tabular data files (spreadsheets or CSV) of up to 100MB. The system provides several ways of visualizing the data (charts, maps, timelines), and the ability to query by filtering and aggregating the data. We support integrating data from multiple sources by performing joins across tables that may belong to different users. Users can keep the data private, share it with a select set of collaborators, or make it public. The discussion feature of Fusion Tables allows collaborators to conduct detailed discussions of the data at the level of individual rows, columns or cells. Given our target audience, we converged on a few principles that underly the design of service. First, our goal is to offer a tool that makes data management easier, and therefore approachable to a larger audience of users. The users we are targeting do not necessarily have any training in using database systems, and typically do not have access to expert DBAs. Second, it was clear that one of the main impediments to data sharing and integration was lack of {\em incentive}. Hence, we built several mechanisms to provide incentives for data sharing. Third, we focused on exploring and supporting features that enable users to {\em collaborative} effectively on data management in the cloud. In particular, when data is shared among multiple collaborators, querying is only one part of the activity, and the system needs to support the process of agreeing on the meaning of data, and discussion on possible errors it may contain. Fusion Tables was inspired by several existing tools. ManyEyes and Swivel enable users to upload data and visualize it in several ways. We go further by providing data management capabilities and a sharing model that does not requiring making your data always public. We strive to preserve the ease of use provided by spreadsheets, but adapt it to larger data sets where the data and the presentation need not necessarily be one of the same. Several online database management tools exist (e.g., DabbleDB, Socrata, Zoho Creator), but Fusion Tables focuses on the collaboration aspects of data management and handles larger data sets. The development of Fusion Tables was carried out under a very tight time schedule, and in addition to the novel features of the system, we also had to support the typical functionality expected from a data management service. The current service represents the features we deemed highest priority for an initial launch. This paper and the accompanying talk describe some of the design choices underlying Fusion Tables, some of the underlying architecture choices we built upon, and some initial feedback from users.

Visual Interfaces to Data

Chris Stolte (Tableau Software)

Easy-to-use visual interfaces to data can broadly expand the audience for databases. Domain experts rather than database experts can engage in rapid-fire Q&A sessions with the data. Visual interfaces can provide a medium for story-telling, debate, and conversations about the data. They can also put new and challenging demands on the capabilities of traditional relational databases. In this talk, I will describe our formal language-based approach to visual analysis and how the use of a formal language enables us to build user experiences that more effectively support the process of analysis. Tableau's VizQL algebra is a declarative language for succinctly describing visual representations of data and analytics operations on the data. A VizQL statement compiles into the SQL or MDX queries necessary to generate the view and into the graphical commands to render the interactive view of the data. Our easy-to-use drag-and-drop user experiences for analysis and visual interface authoring are built on top of VizQL. In addition to supporting the process of analysis, a formal language-based approach provides a basis for reasoning about the structure of views and the space of possible views. This in turn enables the development of powerful new analytic capabilities, such as automatic presentation of structured data, visual authoring of statistical models, and view-based calculation, which we demonstrate. I will also discuss the challenges we have faced in getting relational databases "in the wild" to effectively support visual analysis for the average business or scientific user. The challenges range from the technical to the political. Traditional relational databases, both for OLTP and OLAP, often require sophisticated data modeling and data management expertise, optimize for performance based on known workloads, and are designed for

scaling to large databases sizes (e.g. PB or TB) on clusters of machines rather than reducing analytic latency using limited hardware. I will describe our approaches to building a database focused on providing interactive query performance on tens or hundreds of millions of rows of data with little or no data modeling (physical or logical) and running on a typical knowledge worker desktop machine. Finally, I will discuss the changing landscape of interfaces to databases. The original interface to the database was transactional in focus: Many users read and make atomic changes to a small number of rows in a large database. In recent decades, powerful analytic use cases have emerged focused on the study and analysis of massive amounts of data by relatively small numbers of power users. The emergence of easily authored visual interfaces to public and private data changes will enables a new style of database usage. Millions of users performing analytics on thousands of data sets all hosted in the cloud with usage demonstrating the familiar long-tail distribution. Everyone will become an author and all interfaces will enable analytics.

Graphical XQuery in the AquaLogic Data Services Platform

Vinayak Borkar, Michael Carey (UC Irvine), Sebu Koleth, Alex Kotopoulis (Oracle), Kautul Mehta (SAP), Joshua Spiegel, Sachin Thatte (Oracle), Till Westmann (SAP)

The AquaLogic Data Services Platform (ALDSP) is a middleware platform developed at BEA Systems for building services, referred to as data services, that integrate, access, and manipulate information coming from multiple heterogeneous sources of data (including databases, files, and other services). ALDSP uses functions that produce and consume XML to model heterogeneous information sources, so the integration logic for data access services in ALDSP is specified declaratively using the XQuery language. A challenge that we faced in developing ALDSP was providing effective graphical tooling to help data service developers to develop information integration queries. In this paper, we describe the graphical XQuery Editor (XQE) that resulted from our attempt to tackle this challenge. XQE handles the full XQuery language and provides a robust two-way editing experience involving both graphical and source views of each query. XQE is novel in being the first commercial graphical XQuery editor to support both of these features.

SIGMOD Industrial Session 5: Stream Processing

Chair: Marios Hadjieleftheriou (AT&T Research)

Thursday 10:30 – 12:00

Analytics over Continuous and DisContinuous (ACDC) Streams: The Truviso Approach

Sailesh Krishnamurthy, Rushan Chen, Jeffery Davis, Daniel Farina, Michael Franklin, Alan Li, Neil Thombre (Truviso)

Streaming continuous analytics systems have emerged as key solutions for dealing with massive data volumes and demands for low latency. These systems have been heavily influenced by an assumption that data streams can be viewed as sequences of ordered continuous data. The reality, however, is that streams are not continuous and disruptions of various sorts in the form of either big chunks of late arriving data or arbitrary failures are endemic. We argue, therefore, that stream processing needs a fundamental rethink and advocate a unified approach providing Analytics over Continuous and DisContinuous (ACDC) streams of data. Our approach is based on a simple insight – partially process independent runs of data and defer the consolidation of the associated partial results to when the results are actually used on an on demand basis. Not only does our approach provide the first real solution to the problem of data that arrives arbitrarily late, it also lets us solve a host of hard problems such as parallelism, recovery, transactional consistency, high availability, failover, and replication that have been neglected by streaming systems. In this paper we describe the Truviso ACDC approach and outline some of the key technical arguments and insights behind it.

IBM Infosphere Streams for Scalable, Real-time, Intelligent Transportation Services

Alain Biem, Eric Bouillet, Hanhua Feng, Anand Ranganathan, Anton Ribov, Olivier Verscheure (IBM TJ Watson), Haris Koutsopoulos, Carlos Moran (KTH)

With the widespread adoption of location tracking technologies like GPS, the domain of intelligent transportation services has seen growing interest in the last few years. Services in this domain make

use of real-time location-based data from a variety of sources, combine this data with static location-based data such as maps and points of interest databases, and provide useful information to end-users. Some of the major challenges in this domain include i) scalability, in terms of processing large volumes of real-time and static data; ii) extensibility, in terms of being able to add new kinds of analyses on the data rapidly, and iii) user interaction, in terms of being able to support different kinds of one-time and continuous queries from the end-user. In this paper, we demonstrate the use of IBM Infosphere Streams, a new IBM product and a scalable stream processing platform, for tackling these challenges. We describe a prototype system that generates dynamic, multi-faceted views of transportation information for the city of Stockholm, using real vehicle GPS and road-network data. The system also continuously derives current traffic statistics, and refines traffic forecast models accordingly. Finally the system continuously updates shortest-time routes from real-time observed and inferred traffic conditions. Our performance experiments illustrate the scalability of the system. For instance, our system can process over 250000 incoming GPS points per second, combine it with a map containing over 600,000 links, continuously generate different kinds of traffic statistics and answer user queries.

SIE-OBI: A Streaming Information Extraction Platform for Operational Business Intelligence

Malu Castellanos, Chetan Gupta, Umesh Dayal, Song Wang (HP Labs)

Today when so much data is available at our fingertips, we fall rather short in our capability to consume it at the pace that it is generated. Enterprises today not only have mountains of internal data but also endless external data from the Web can even be available as soon as it is generated (e.g., RSS feeds). Social media like news feeds, tweets and blogs, among others, provide the opportunity to get informed instantly, but only if we are able to digest these oceans of data. In the context of an enterprise, the capability to extract valuable information from all sorts of data (i.e., structured, semi-structured and unstructured) and act instantly either reactively or proactively, provides a tremendous competitive advantage. In particular, identifying external events that may affect the enterprise operations and objectives is essential in the intelligent enterprise. Typically, enterprise data warehouses (EDW) serve the vital function of being a single version of the truth for historical business reporting but they fall short when it comes to providing a real-time view for operational decision making. Moreover, in general EDWs are fed from structured sources like relational databases and ignore unstructured sources like document repositories and social media feeds. As a consequence, the opportunity to gain competitive advantage from the capability of correlating unstructured historical data with the unstructured streaming data from the web in a timely manner is lost. This is the case with large enterprises that have thousands of customers and partners all over the world and a myriad of legacy contracts of a variety of types with them. Data buried in the legalese of contracts is not being used to make business decisions upon the occurrence of world events that may affect contractual relationships. For example: political instability in a country (what contracts exist with suppliers based in this country?), significant fluctuations in currency values (what contracts are denominated in this currency?), changes in commercial law (how does the change affect the risk in each contract?), mergers and acquisitions (what contracts exist with the parties involved in the merger?). To our surprise, we found out in our investigations that this is a common reality that extends to many different domains in industry. With SEI-OBI, we address this problem.

SIGMOD DEMONSTRATION ABSTRACTS

SIGMOD Demonstrations Session Group A: Cloud, OLAP, and XML
Tuesday 10:30 – 12:00 and Wednesday 10:30 – 12:00

HadoopDB in Action: Building Real World Applications

Kamil Bajda-Pawlikowski, Azza Abouzeid, Jiewen Huang, Daniel Abadi, Avi Silberschatz (Yale University)

HadoopDB is a hybrid of MapReduce and DBMS technologies, designed to meet the growing demand of analyzing massive datasets on very large clusters of machines. Our previous work has shown that HadoopDB approaches parallel databases in performance and still yields the scalability and fault tolerance of MapReduce-based systems. In this demonstration, we focus on HadoopDB's flexible architecture and versatility with two real world application scenarios, namely, a semantic web data application for protein sequence analysis and a business data warehousing application based on TPC-H. The demonstration offers a thorough walk-through of how to easily build applications on top of HadoopDB.

Online Aggregation and Continuous Query support in MapReduce

Tyson Condie, Neil Conway, Joseph Hellerstein, Peter Alvaro (UC Berkeley), John Gerth, Justin Talbot (Stanford University), Khaled Elmeleegy, Russell Sears (Yahoo! Research)

MapReduce is a popular framework for data-intensive distributed computing of batch jobs. To simplify fault tolerance, the output of each MapReduce task and job is materialized to disk before it is consumed. During this demonstration, we describe a modified MapReduce architecture that allows data to be pipelined between operators. This extends the MapReduce programming model beyond batch processing, and can reduce completion times and improve system utilization for batch jobs as well. We demonstrate a modified version of the Hadoop MapReduce framework that supports online aggregation, which allows users to see ``early returns" from a job as it is being computed. Our Hadoop Online Prototype (HOP) also supports continuous queries, which enable MapReduce programs to be written for applications such as event monitoring and stream processing. HOP retains the fault tolerance properties of Hadoop, and can run unmodified user-defined MapReduce programs.

MapDupReducer: Detecting Near Duplicates over Massive Datasets

Chaokun Wang, Jianmin Wang (Tsinghua University), Xuemin Lin, Wei Wang (University of New South Wales), Haixun Wang, Hongsong Li, (Microsoft Research, Asia), Wanpeng Tian, Jun Xu, Rui Li (Tsinghua University)

Near duplicate detection benefits many applications, e.g., on-line news selection over the Web by keyword search. The purpose of this demo is to show the design and implementation of MapDupReducer, a MapReduce based system capable of detecting near duplicates over massive datasets efficiently.

Large Graph Processing in the Cloud

rishan Chen, Xuetian Weng, Bingsheng He, Mao Yang (MSRA), Bo Peng (PKU)

As the study of graphs, such as web and social graphs, becomes increasingly popular, the requirements of efficiency and programming flexibility of large graph processing tasks challenges existing tools. We propose to demonstrate Surfer, a large graph processing engine designed to execute in the cloud. Surfer provides two basic primitives for programmers – MapReduce and propagation. MapReduce, originally developed by Google, processes different key-value pairs in parallel, and propagation is an iterative process that transfers information along each edge from a vertex to its neighbors in the graph. These two primitives are complementary in graph processing. MapReduce is suitable for processing flat data structures, such as vertex-oriented tasks, and propagation is optimized for edge-oriented tasks on partitioned graphs. Surfer consists of a small set of high level building blocks that use these two primitives. Developers may also construct custom

building blocks. Surfer further provides a GUI (Graphical User Interface) using which developers can visually create large graph processing tasks. Surfer transforms a task into an execution plan composed of MapReduce and propagation operations. It then automatically applies various optimizations to improve the efficiency of distributed execution. Surfer also provides a visualization tool to monitor the detailed execution dynamics of the execution plan to show the interesting tradeoffs between MapReduce and propagation. We demonstrate our system in two ways: first, we demo the ease-of-programming features of the system; second, we show the efficiency of the system with a series of applications on a real large graph representing users and their social connections in the MSN social network. This graph has 500 million vertices and 30 billion undirected edges with a total memory footprint of 120GB. We find that Surfer is simple to use and is highly efficient for large graph-based tasks.

DCUBE: Discrimination Discovery in Databases

Salvatore Ruggieri, Dino Pedreschi, Franco Turini, (Universita di Pisa)

Discrimination discovery in databases consists in finding unfair practices against minorities which are hidden in a dataset of historical decisions. The DCUBE system implements the approach of [1], which is based on classification rule extraction and analysis, by centering the analysis phase around an Oracle database. The proposed demonstration guides the audience through the legal issues about discrimination hidden in data, and through the legally-grounded analyses to unveil discriminatory situations. The SIGMOD attendees will freely pose complex discrimination analysis queries over the database of extracted classification rules, once they are presented with the database relational schema, a few ad-hoc functions and procedures, and several snippets of SQL queries for discrimination discovery. [1] D. Pedreschi, S. Ruggieri, and F. Turini. Discrimination-aware data mining. In Proc. of KDD 2008, pages 560--568, ACM, 2008. Extended version to appear in ACM Trans. on Know. Disc. from Data.

S-OLAP: An OLAP system for analyzing sequence data

Chun Kit Chui, Ben Kao (University of Hong Kong), Eric Lo (Hong Kong Polytechnic University), David Cheung (University of Hong Kong)

The Sequence OLAP (S-OLAP) system is a novel online analytical processing system for analyzing sequence data. S-OLAP supports “pattern-based” grouping and aggregation on sequence data — a very powerful concept and capability that is not supported by traditional OLAP systems. It also supports several new OLAP operations that are specific to sequence data analysis. The query processing techniques documented in previous research works have been implemented in our S-OLAP engine for efficient query processing. The system also provides users with a friendly graphical interface for query construction and result visualization. Query parameters can be interactively refined and the results are updated in real-time so as to facilitate the exploratory analysis of sequence data.

ProgXe: Progressive Result Generation Framework for Multi-Criteria Decision Support Queries

Venkatesh Raghavan, Elke Rundensteiner (Worcester Polytechnic Institute)

We demonstrate ProgXe, a practical approach to support Multi-Criteria Decision Support (MCDS) applications that need to report results as they are being generated to enable the user to make competitive decisions. ProgXe transforms the execution of MCDS queries involving skyline over joins to be non-blocking, by progressively generating results early and often. The demo highlights key feature of our progressive execution framework that optimizes for early output generation by: (1) evaluating the query at multiple levels of abstraction, (2) exploiting the skyline knowledge gained from both input as well as mapped output spaces. The audience will be able to submit MCDS queries over two real-world applications. We provide visualization tools that enable the user to make quick decisions, compare alternative techniques, and provide capability to fine-tune the query predicates based on the early output results.

XTaGe: A Flexible XML Collection Generator

Maria Perez, Ismael Sanz, Rafael Berlanga (Universitat Jaume I)

In this demonstration we present XTaGe (XML Tester and Generator), a flexible system for the creation of complex XML collections. XTaGe focuses on XML collections with complex structural constraints and domain-specific characteristics, which would be very difficult or impossible to replicate using existing XML generators. It addresses the limitations of current XML generators by providing a highly extensible framework to introduce controlled variability in XML structures.

K*SQL: A Unifying Engine for Sequence Patterns and XML
Barzan Mozafari, Kai Zeng, Carlo Zaniolo (UCLA)

A strong interest is emerging in SQL extensions for sequence patterns using Kleene-closure expressions. This burst of interest from both the research community and the commercial world is due to the many database and data stream applications made possible by these extensions-including financial services, RFID-based inven- tory management, and electronic health systems. In this demo we will present the K*SQL system that represents a major step for- ward in this area. In fact, K*SQL supports generalized Kleene- closure queries and also achieves the expressive power of the nested word model. This greatly expands SQL application domain to include XML queries, software trace analysis, and genomics. In this demo, we will first demonstrate the optimized pattern search techniques on which our efficient K*SQL engine is based, and its use in supporting complex pattern queries expressed in SQL. Then, we will demonstrate how Xpath queries on streaming XML documents are compiled, optimized and executed efficiently on our system as K*SQL queries. We will finally demonstrate simple software trace analysis and genomic applications supported in our K*SQL system.

SIGMOD Demonstrations Session Group B: Stream, keyword search, and Web
Tuesday 13:30 – 15:00 and Wednesday 16:00 – 17:30

Symbiote - A Reconfigurable Logic Assisted Data Stream Management System (RLADSMS)
Pranav Vaidya, Jaehwan John Lee, Fracis Bowen, Yingzi Du, Chandima Hewa Nadungodage, Yuni Xia (IUPUI)

Numerous monitoring applications such as traffic control systems, border patrol monitoring, and person locator services generate a large number of multimedia data streams that need to be analyzed and processed using image processing and data stream management techniques in order to detect significant events of interest or abnormal conditions. Such multimedia monitoring systems usually have high-bandwidth characteristics, stricter real-time deadlines, and high accuracy requirements. In an attempt to meet all of these requirements, we have designed Symbiote - a distributed Reconfigurable Logic Assisted multimedia Data Stream Management System (RLADSMS) at Indiana University Purdue University, Indianapolis (IUPUI) that provides hardware accelerated data stream processing using Field Programmable Gate Arrays (FPGAs).

Interactive Visual Exploration of Neighbor-Based Patterns in Data Streams
Di Yang, Zhenyu Guo, Zaixian Xie, Elke Rundensteiner, Matthew Ward (Worcester Polytechnic Institute)

We will demonstrate our system, called V iStream, support- ing interactive visual exploration of neighbor-based patterns [7] in data streams. V istream does not only apply innovative multi-query strategies to compute a broad range of popular patterns, such as clusters and outliers, in a highly efficient manner, but it also provides a rich set of visual interfaces and interactions to enable real-time pattern ex- ploration. With ViStream, analysts can easily interact with the pattern mining processes by navigating along the time horizons, abstraction levels and parameter spaces, and thus better understand the phenomena of interest.

TwitterMonitor: Trend Detection over the Twitter Stream
Michael Mathioudakis, Nick Koudas (University of Toronto)

We present TwitterMonitor, a system that performs trend detection over the Twitter stream. The

system identifies emerging topics (i.e. `trends') on Twitter in real time and provides meaningful analytics that synthesize an accurate description of each topic. Users interact with the system by ordering the identified trends using different criteria and submitting their own description for each trend. We discuss the motivation for trend detection over social media streams and the challenges that lie therein. We then describe our approach for trend detection, as well as the architecture of TwitterMonitor. Finally, we lay out our demonstration scenario.

Glacier: A Query-to-Hardware Compiler

Rene Mueller, Jens Teubner, Gustavo Alonso (ETH, Zurich)

Field-programmable gate arrays (FPGAs) are a promising technology that can be used in database systems. In this demonstration we show Glacier, a library and a compiler that can be employed to implement streaming queries as hardware circuits on FPGAs. Glacier consists of a library of compositional hardware modules that represent stream processing operators. Given a query execution plan, the compiler instantiates the corresponding components and wires them up to a digital circuit. The goal of this demo is to show the flexibility of the compositional approach.

Exploratory Keyword Search on Data Graphs

Hilit Achiezra, Konstantin Golenberg (The Hebrew University), Benny Kimelfeld (IBM Almaden), Yehoshua Sagiv (The Hebrew University)

A demo of a system for keyword search on data graphs is proposed. The system supports search, exploration and question answering. The demo is on two challenging datasets, the full DBLP (enhanced with abstracts of some of the publications) and Mondial (which is highly cyclic and has a complex schema). The search algorithm employed by the system has important, provable properties: polynomial delay, 2-approximate ranking by increasing height, and ability to find all the answers. The system uses a novel technique for displaying multi-node subtrees in a compact graphical form that facilitates quick and easy understanding, which is essential for effective browsing of the answers.

Integrating Keyword Search with Multiple Dimension Tree Views over a Summary Corpus Data Cube

Mark Sifer (University of Wollongong), Yutaka Watanobe, Subhash Bhalla (University of Aizu, Japan)

We demonstrate a system that integrates a novel OLAP component with a keyword search engine, to support querying over sparse and ragged corpus data. The key contribution of our system is the integration of dynamically selected point sets such as search results with OLAP querying over aggregated data. During the demonstration, participants will be able to enter a keyword search; observe the returned list of result files; observe distributional features such as outliers and clusters of results in the corpus in multiple dimension views; and select and partition corpus slices in the OLAP component to narrow search results. Participants will be able to experience not just the individual querying features of our system, but the way that they work together to facilitate smooth interaction sequences that combine OLAP and keyword search querying.

Query Portals: Dynamically Generating Portals for Web Search Queries

Sanjay Agrawal, Kaushik Chakrabarti, Surajit Chaudhuri, Venkatesh Ganti, Christian Konig, Dong Xin (Microsoft Research)

Many web search queries seek information about entities in structured databases, say, about movies, people, and products. For example, many product related queries on search engines (Amazon, Google, Yahoo, Live Search) are answered by searching underlying product databases containing names, descriptions, specifications, and reviews of products. However, current vertical search engines are silo-ed in that their results are independent of those from web search. This often leads to empty or incomplete results, as query terms are matched against the information in the underlying database only. In order to overcome this problem, we propose an approach that first identifies relationships between web documents and entities in structured databases. We leverage these relationships to surface for each search query highly relevant entities from a given structured

database. We demonstrate that the combination of search results and structured data creates a rich set of results, for the user to focus on and refine their search.

Creating and Exploring Web Form Repositories

Luciano Barbosa, Hoa Nguyen, Nguyen Thanh, Ramesh Pinnamaneni, Juliana Freire (University of Utah)

We present DeepPeep (http://www.deeppeep.org), a new system for discovering, organizing and analyzing Web forms. DeepPeep allows users to explore the entry points to hidden-Web sites whose contents are out of reach for traditional search engines. Besides demonstrating important features of DeepPeep and describing the infrastructure we used to build the system, we will show how this infrastructure can be used to create form collections and form search engines for different domains. We also present the analysis component of DeepPeep which allows users to explore and visualize information in form repositories, helping them not only to better search and understand forms in different domains, but also to refine the form gathering process.

SIGMOD Demonstrations Session Group C: Schema, language, and spatial
Tuesday 15:30 – 17:00 and Thursday 10:30 – 12:00

Exploring Schema Similarity At Multiple Resolutions

Ken Smith, Craig Bonaceto, Chris Wolf, Beth Yost, Michael Morse, Peter Mork, Doug Burdick (MITRE)

Large, dynamic, and ad-hoc organizations must frequently initiate integration and data sharing efforts without sufficient awareness of data source relationships. We demonstrate an integrated environment for exploring schema similarity relationships across multiple resolutions. Its centerpiece is Affinity, a novel tool for displaying and interacting with clusters of related schemas in a schema repository. Affinity allows drill-down to examine the extent and content of schema overlap within a cluster. Further drill down allows users to examine fine-grained element-level correspondences between between selected schemas, via a seamlessly integrated schema matching tool.

An Automated, yet Interactive and Portable DB designer

Ioannis Alagiannis (EPFL) Debabrata Dash (EPFL & Carnegie Mellon University), Karl Schnaitter (UC Santa Cruz), Anastasia Ailamaki (EPFL), Neoklis Polyzotis (UC Santa Cruz)

Tuning tools attempt to configure a database to achieve optimal performance for a given workload. Selecting an optimal set of physical structures is computationally hard since it involves searching a vast space of possible configurations. Commercial DBMS offer tools that can address this problem. Their usefulness, however, is limited by their dependence on greedy heuristics, requirement of an offline workload, and lack of an optimal materialization schedule to get the best out of suggested design features. Moreover, the open source DBMS do not provide any automated tuning tools. This demonstration introduces a comprehensive physical designer for the PostgreSQL open source DBMS. The tool suggests design features for both offline and online workloads. It provides close to optimal suggestions for indexes for a given workload by modeling the problem as a combinatorial optimization problem and solving it by sophisticated and mature solvers. It also determines the interaction between indexes to suggest optimal materialization strategy for the selected indexes. The tool is interactive as it allows the database administrator (DBA) to suggest a set of candidate features and shows their benefits and interactions visually. We demonstrate the usefulness of the tool using large real-world scientific datasets and query workloads.

Midas: Integrating Public Financial Data

Sreeram Balakrishnan (IBM - Silicon Valley Lab), Vivian Chu, Mauricio Hernandez, Howard Ho (IBM Research – Almaden), Rajasekar Krishnamurthy, Liu Shi (IBM Research – China), Jan Pieper (IBM Research – Almaden), Jeffrey Pierce, Lucian Popa, Christine Robson (IBM Research – Almaden), Lei Shi, (IBM Research – China), Ioana Stanoi (IBM Research – Almaden), Edison Ting (IBM - Silicon Valley Lab), Shivakumar Vaithyanathan, Huahai Yang (IBM Research – Almaden)

We present Midas, a system that unleashes the value of information buried in regulatory filing submitted to the U.S. Securities and Exchange Commission (SEC) by extracting, conceptualizing, integrating, and aggregating data from semi-structured or text form. Midas was built from the ground up this year and is a scalable Hadoop-based system currently running at IBM Research - Almaden. The major differentiators from existing products/services are the automatic ability to extract concepts, attributes and relationships from text, and enabling complex analysis over this extracted information. In this demo, we will showcase the 25,000 merged and cleansed financial company and people objects automatically extracted from SEC data. We will also discuss Midas' architecture and the techniques used to extract and integrate the data.

Worry-Free Database Upgrades: Automated Model-Driven Evolution of Schemas and Complex Mappings

James Terwilliger, Philip Bernstein, Adi Unnithan (Microsoft Corporation)

Schema evolution is an unavoidable consequence of the application development lifecycle. The two primary schemas in an application, the client object model and the persistent database model, must co-evolve or risk quality, stability, and maintainability issues. We present MoDEF, an extension to Visual Studio that supports automatic evolution of object-relational mapping artifacts in the Microsoft Entity Framework. When starting with a valid mapping between client and store, MoDEF translates changes made to a client model into incremental changes to the store as an upgrade script, along with a new valid mapping to the new store. MoDEF mines the existing mapping for mapping patterns which MoDEF reuses for new client artifacts.

US-SQL: Managing Uncertain Schemata

Matteo Magnani, Danilo Montesi (University of Bologna)

In this paper we propose a demonstration on the management of uncertain schemata. Many works have studied the problem of representing uncertainty on attribute values or tuples, like the fact that a value is 10 with probability .3 or 20 with probability .7, leading to the implementation of probabilistic database management systems. In our demo we deal with the representation of uncertainty about the metadata, i.e., about the meaning of these values. The audience of the demo will be able to create alternative probabilistic schemata on a database, execute queries over uncertain schemata and verify how this additional information is stored in an underlying relational database and how queries are executed.

PAROS: Pareto Optimal Route Selection

Matthias Schubert, Franz Graf, Matthias Renz, Hans-Peter Kriegel (Ludwig-Maximilians-Universitaet Munchen)

Modern maps provide a variety of information about roads and their surrounding landscape allowing navigation systems to go beyond simple shortest path computation. In this demo, we show how the concept of skyline queries can be successfully adapted to routing problems considering multiple road attributes. In particular, we demonstrate how to compute several pareto-optimal paths which contain optimal results for a variety of user preferences. The PAROS-system has two main purposes. The first is to calculate the route skyline for a starting point and a destination. Our demonstrator visualizes the result set for up to three road attributes. Therefore, we provide a dual view on the computed skyline paths. The first view displays the result paths on the road map itself. The second view describes the result paths in the property space, displaying the trade-off between the underlying criteria. Thus, a user can browse through the results in order to find the path which fits best to his personal preferences. The second component of our system suits analysis issues. In this component, we illustrate the functionality of the underlying route skyline algorithm. Thus, we provide benchmark information about processing time and the search space visited during route skyline computation.

MoveMine: Mining Moving Object Databases

Zhenhui Li (UIUC), Jae-Gil Lee (IBM Almaden Research Center), Jiawei Han (UIUC)

With the maturity of GPS, wireless, and Web technologies, increasing amounts of movement data

collected from various moving objects such as animals, vehicles, mobile devices, and climate radars have become widely available. Analyzing such data has broad applications in ecological study, vehicle control, mobile communication management, and climatological forecast. However, few data mining tools are available for flexible and scalable analysis of massive-scale moving object data. Our system, MoveMine, is designed for sophisticated moving object data mining by integrating several attractive functions including moving object pattern mining and trajectory mining. We use state-of-the-art and novel techniques in implementation of the functions. A user-friendly interface will be provided to facilitate interactive exploration of mining results and flexibly tune the underlying methods. Since Movement is tested on several different real data sets, it will benefit moving object data owners to carry out various kinds of analysis on their data. At the same time, it will benefit database researchers to realize the importance and limitations of current techniques as well as the potential future studies in moving object data mining.

PIQL: A Performance Insightful Query Language

Michael Armbrust, Stephen Tu, Armando Fox, Michael Franklin, David Patterson, Nick Lanham, Beth Trushkowsky, Jesse Trutna (UC Berkeley)

Large-scale websites are increasingly moving from relational databases to distributed key-value stores for high request rate, low latency workloads. Often this move is motivated not only by key-value stores ability scale simply by adding more hardware, but also by the easy to understand predictable performance they provide for all operations. While this datamodel works well for system when lookups are only done by primary key, such as the shopping cart example from Dynamo, more complex queries require onerous explicit index management and imperative data lookup by the developer. We propose PIQL, a Performance Insightful Query Language that allows developers to express many of the queries found on these websites, while still providing strict performance guarantees.

DoCQS: A Prototype System for Supporting Data-oriented Content Query

Mianwei Zhou, Tao Cheng, Kevin Chang (University of Illinois at Urbana-champaign)

Witnessing the richness of data in document content and many ad-hoc efforts for finding such data, we propose a Data-oriented Content Query System(DoCQS), which is oriented towards fine granularity data of all types by searching directly into document content. DoCQS uses the relational model as the underlying data model, and offers a powerful and flexible Content Query Language(CQL) to adapt to diverse query demands. In this demonstration, we show how to model various search tasks by CQL statements, and how the system architecture efficiently supports the CQL execution. Our online demo of the system is available at http://falcon4.cs.uiuc.edu:8080/DoCQS.

SIGMOD Demonstrations Session Group D: New technology, and potpourri
Wednesday 8:30 – 10:00 and Thursday 13:30 – 15:00

QRelX: Generating Meaningful Queries that Provide Cardinality Assurance

Manasi Vartak, Venkatesh Raghavan, Elke Rundensteiner (Worcester Polytechnic Institute)

In many business and consumer applications, queries have cardinality constraints. However, since there is minimal support for cardinality assurance, users must follow a trial-and-error approach to find queries with the desired result size. In this demonstration, we present QRelX -- a novel framework to automatically generate alternate queries that maintain closeness to the original user query and satisfying cardinality constraints. Our demonstration is an interactive game that illustrates the importance of cardinality assurance as well as our proposed solution. The game allows the audience to compete with QRelX via manual query refinement and provides a visual representation of the innovative strategies used by QRelX.

Performing Sound Flash Device Measurements: The uFLIP Experience

Matias Bjorling (University of Copenhagen), Lionel Le Folgoc, Ahmed Mseddi (INRIA), Philippe Bonnet (University of Copenhagen), Luc Bouganim (INRIA), Bjorn Jonsson (Reykjavik University)

It is amazingly easy to get meaningless results when measuring flash devices, partly because of the peculiarity of flash memory, but primarily because their behavior is determined by layers of complex, proprietary, and undocumented software and hardware. In this demonstration, we share the lessons we learned developing the uFLIP benchmark and conducting experiments with a wide range of flash devices. We illustrate the problems that are actual obstacles to sound performance and energy measurements, and we show how to mitigate the effects of these problems. We also present the uFLIP web site and its on-line visualization tool that should help the research community investigate flash device behavior.

GDR: A System for Guided Data Repair

Mohamed Yakout, Ahmed Elmagarmid, Jennifer Neville, Mourad Ouzzani (Purdue University)

Improving data quality is a time-consuming, labor-intensive and often domain specific operation. Existing data repair approaches are either fully automated or not efficient in interactively involving the users. We present a demo of GDR, a Guided Data Repair system that uses a novel approach to efficiently involve the user alongside automatic data repair techniques to reach better data quality as quickly as possible. Specifically, GDR generates data repairs and acquire feedback on them that would be most beneficial in improving the data quality. GDR quantifies the data quality benefit of generated repairs by combining mechanisms from decision theory and active learning. Based on these benefit scores, groups of repairs are ranked and displayed to the user. User feedback is used to train a machine learning component to eventually replace the user in deciding on the validity of a suggested repair. We describe how the generated repairs are ranked and displayed to the user in a "usefullooking" way and demonstrate how data quality can be effectively improved with minimal feedback from the user.

Crescando

Georgios Giannikis, Philipp Unterbrunner (ETH Zurich), Jeremy Meyer (Amadeus), Gustavo Alonso (ETH Zurich), Dietmar Fauser (Amadeus), Donald Kossman (ETH Zurich)

This demonstration presents Crescando, an implementation of a distributed relational table that guarantees predictable response time on unpredictable workloads. In Crescando, data is stored in main memory and accessed via full-table scans. By using scans instead of index lookups, Crescando overcomes the read-write contention in index structures and eliminates the scalability issues that exist in traditional in\-dex-based systems. Crescando is specifically designed to process a large number of queries in parallel, allowing high query rates. The goal of this demonstration is to show the ability of Crescando to a) quickly answer arbitrary user-generated queries, and b) execute a large number of queries and updates in parallel, while providing strict response time and data freshness guarantees.

A Tool for Configuring and Visualizing Database Parameters

Vamsidhar Thummala, Shivnath Babu (Duke University)

We present iTuned, a tool for configuring database parameters. iTuned takes a SQL workload as input and recommends good settings for configuration parameters. In addition, iTuned provides response surface and sensitivity analysis plots that will help the end user analyze the the impact of each parameter. iTuned has the following novel features: (i) a technique called Adaptive Sampling that pro actively brings in appropriate data through planned experiments to find high-impact parameters and high-performance parameter settings, (ii) an executor that supports on-line experiments in production database environments through a cycle-stealing paradigm that places near-zero overhead on the production workload; and (iii) portability across different database systems. In this demonstration, we focus on the interesting use case scenarios of iTuned and show the effectiveness of recommended configuration settings on standard database benchmarks.

Pluggable Personal Data Servers

Nicolas Anciaux, Luc Bouganim, Yanli Guo, Philippe Pucheral (INRIA), jean-jacques vandewalle (Gemalt), Shaoyi Yin (INRIA)

The amount of personal information collected by governments, corporations, commercial Web sites, public and private agencies, is increasing at a high rate. The suspicion of individuals towards this ineluctable centralization of personal data is fueled by computer security surveys pointing out the vulnerability of database servers against external and internal attacks. Personal databases can be hosted in secure hardware but this solution applied so far only to small mono-user databases. Today, a new generation of secure hardware devices emerges and drastically transforms the way personal data can be managed. Secure tokens combine the hardware security of smart card microcontrollers, the storage capacity of memory sticks and the performance and universality of the USB communication protocol. Currently, a secure token is primarily seen as a raw secure repository, i.e., a set of documents protected by the tamper-resistance of the token and unlocked on demand by the record owner thanks to a PIN code. This demonstration paper advocates a much broader exploitation of the secure token storage and computing capabilities. It suggests the idea of embedding in secure hardware the complete chain of software usually found on traditional servers, that is a Web server, a set of servlet-based applications and a DBMS engine managing the on-board database and enforcing powerful access control rules. The resulting system, named PlugDB, builds upon a set of research results published earlier. The demonstration scenario itself is based on a real experiment of secure and mobile healthcare folders used to improve care coordination at home for dependent people. Beyond this experiment, we expect that PlugDB will contribute to open new ways of thinking about and organizing the management of personal data.

Mask: A System for Privacy-Preserving Policy-Based Access to Published Content

Mohamed Nabeel, Ning Shang, John Zage, Elisa Bertino (Purdue University)

We propose to demonstrate Mask, the first system addressing the seemingly-unsolvable problem of how to selectively share contents among a group of users based on access control policies expressed as conditions against the identity attributes of these users while at the same time assuring the privacy of these identity attributes from the content publisher. Mask consists of three entities: a Document Publisher, Users referred to as Subscribers, and Identity Providers that issue certified identity attributes. The document publisher specifies access control policies against identity attributes of subscribers indicating which conditions the identity attributes of a subscriber must verify in order for this subscriber to access a document or a subdocument. The main novelty of Mask is that, even though the publisher is able to match the identity attributes of the subscribers against its own access control policies, the publisher does not learn the values of the identity attributes of the subscribers; the privacy of the authorized subscribers is thus preserved. Based on the specified access control policies, documents are divided into subdocuments and the subdocuments having different access control policies are encrypted with different keys. Subscribers derive the keys corresponding to the subdocuments they are authorized to access. Key distribution in Mask is supported by a novel group key management protocol by which subscribers can reconstruct the decryption keys from the subscription information they receive from the publisher. The publisher however does not learn which decryption keys each subscriber is able to reconstruct. In this demonstration, we show our system using a healthcare scenario.

SimDB: A Similarity-aware Database System

Yasin Silva, Walid Aref (Purdue University), Paul Larson (Microsoft Research)

The identification and processing of similarities in the data play a key role in multiple application scenarios. Several types of similarity-aware operations have been studied in the literature. However, in most of the previous work, similarity-aware operations are studied in isolation from other regular or similarity-aware operations. Furthermore, most of the previous research in the area considers a standalone implementation, i.e., without any integration with a database system. In this demonstration we present SimDB, a similarity-aware database management system. SimDB supports multiple similarity-aware operations as first-class database operators. We describe the architectural changes to implement the similarity-aware operators. In particular, we present the way conventional operators' implementation machinery is used to support similarity-ware operators. We also show how these operators interact with other similarity-aware and regular operators. In particular, we show the effectiveness of multiple equivalence rules that can be used to extend cost-based query

optimization to the case of similarity-ware operations. SimDB is an open source framework that can be used by other researchers and developers to test and integrate new or improved similarity-ware operators and optimization techniques.

A Demonstration of FlexPref: Extensible Preference Evaluation Inside the DBMS Engine

Justin Levandoski, Mohamed F. Mokbel, Mohamed Khalefa, Venkateshwar Korukanti (University of Minnesota)

This demonstration presents FlexPref, a framework implemented inside the DBMS query processor that enables efficient and extensible preference query processing. Numerous preference evaluation methods for databases have been proposed (e.g., skyline, top-k, k-dominance, k-frequency). Currently, the implementation of these methods at the core of a DBMS is (1) beneficial, as the preference method realizes efficient query processing for arbitrary database queries, but (2) not practical, as each existing (and future) preference method requires a custom query processor implementation. FlexPref addresses this paradox by supporting a wide-array of preference evaluation methods in a single extensible code base. Integration with FlexPref is simple, involving the registration of only three functions that capture the essence of the preference method. Once integrated, the preference method "lives" at the core of the database, enabling the efficient execution of preference queries involving common database operations (e.g, selection, projection, join). We demonstrate the use of FlexPref, implemented inside the query processor of PostgreSQL, by walking through the implementation and use of five state-of-the-art preference methods in a real application scenario.

SIGMOD Programming Contest
Tuesday 13:30 – 15:00, and Wednesday 10:30 – 12:00

At this year's conference, student teams from degree-granting institutions were invited to compete in a programming contest to develop a distributed query engine, built on top on the main-memory indexing system that was the topic of last year's contest. The winning team will be announced and awarded a prize of $5,000 at the conference. Submissions were judged based on their overall performance on a supplied workload. For more information on the contest see: http://dbweb.enst.fr/events/sigmod10contest/.

Five finalist teams have been selected out of 29 competitors. Finalists come from Stanford University (USA), TU Kaiserslautern (Germany), INSA Lyon (France), KAUST (Saudi Arabia), and Saint-Petersburg University (Russia). They present their designs at Demo Session B, on Tuesday 13:30 – 15:00 and Wednesday 10:30 – 12:00. The winner will be announced during the Business Meeting during lunch on Wednesday.

The contest is supported by NSF, Microsoft (platinum sponsors), Amazon, INRIA Saclay (gold sponsors), Exalead, and Yahoo! (silver sponsors).

Organizers: Pierre Senellart (pierre.senellart@telecom-paristech.fr), Clément Genzmer (clement.genzmer@gmail.com).

SIGMOD Undergrad Posters Competition
Monday 18:30 – 21:00

Parallel Algorithms for Top-k Query Processing
Zhenqiang Gong
University of Science and Technology of China (USTC), China

Top-k Algorithms for Interactive Processes
Yael Amsterdamer
Tel Aviv University, Israel

Analyzing Plan Spaces of Query Optimizers
Ian Lorne Charlesworth
University of Waterloo, Canada

Query Optimization on Cloud
Bharath Vissapragada
International Institute of Information Technology, Hyderabad (IIIT-H), India

Utilizing the Quoting System of Online Web Forums to Estimate User Agreement
Emmanouil Karvounis
National and Kapodistrian University of Athens, Greece

Extracting Topics of Debate between Users on Web Discussion Boards
Thodoris Georgiou
National and Kapodistrian University of Athens, Greece

* Undergradute student grants are supported by Google

CONFERENCE CO-LOCATED WORKSHOPS

13th International Workshop on Web and Databases (WebDB 2010)
http://webdb2010.org/
June 6, 2010
Location: Cosmopolitan C

WebDB provides a forum in which researchers, theoreticians, and practitioners can share their insights on problems and solutions at the intersection of data management and the Web. WebDB has high impact and has been a forum in which a number of seminal papers have been presented. This year's WebDB continues in the spirit of previous years.

The themes of this year's WebDB are **Quality of Web Data** and **Linked Data**.

Second International Workshop on Keyword Search on Structured Data (KEYS 2010)
http://keys2010.uni-bonn.de/
June 6, 2010
Location: Cosmopolitan D

Information search is an indispensable component of our lives. Web search engines are widely used for searching textual documents, images, and video. There are also vast collections of structured and semi-structured data both on the Web and in enterprises, such as relational databases, XML, and data extracted from text documents. Traditionally, to access these resources, a user must learn structured query languages, and must be able to access data schemata, which are most likely heterogeneous, complex, and subject to constant change. To relieve web and scientific users from the learning curve and enable them to easily access structured and semi-structured data, there is a growing research interest on supporting keyword search on these data sources. Indeed, the rich meta-information embedded in (semi-)structured data provides a great potential for enhancing search quality compared to search on unstructured documents. However, due to the inherent ambiguity of keyword search and the possibly high complexity of the data, providing intelligent search results efficiently is extremely challenging. The aim of this workshop is to provide a forum for researchers from both academia and industry to discuss the opportunities and challenges in keyword search on (semi-)structured data, and to present the key issues and novel techniques in this area.

Ninth International ACM Workshop on Data Engineering for Wireless and Mobile Access (MobiDE 2010)
http://www.softnet.tuc.gr/mobide10/
June 6, 2010
Location: Discovery A

MobiDE 2010 is the ninth of a successful series of workshops that aims to act as a bridge between the data management, wireless networking, and mobile computing communities. MobiDE has been co-located with the annual SIGMOD conference since 2005. MobiDE 2009 marked the 10-year anniversary of the workshop. This year, MobiDE 2010 is again sponsored by ACM SIGMOD and held in co-operation with ACM SIGMOBILE (pending approval).

First International Workshop on Workflow Approaches to New Data-centric Science (WANDS 2010)

http://wands2010.doc.ic.ac.uk/
June 6, 2010
Location: Discovery B

A number of innovative, but uncoordinated, efforts in data-centric workflows have made their mark on the scientific and business world in recent years. The goal of this workshop is to use these efforts to start bringing together a research community around the theoretical foundations, technology development, and domain applications of workflow systems.

As a first postulate, we take that there is no uniform workflow solution, in the same way that there is no uniform programming language. Furthermore, our goal is not to promote a particular technology, or even interoperability between individual technologies. WANDS 2010 aims to present, discuss, and ultimately better the different computational models, usability patterns, and domain approaches used in the field. The scope of WANDS 2010 includes contributions from both researchers and practitioners on all aspects of data and process management that contribute towards this goal.

Sixth International Workshop on Data Management on New Hardware (DaMoN 2010)

http://event.cwi.nl/damon2010/
June 7, 2010
Location: Regency E

The continued evolution of computing hardware and infrastructure imposes new challenges and bottlenecks to program performance. As a result, traditional database architectures that focus solely on I/O optimization increasingly fail to utilize hardware resources efficiently. CPUs with superscalar out-of-order execution, simultaneous multi-threading, multi-level memory hierarchies, and future storage hardware (such as MEMS) impose a great challenge to optimizing database performance. Consequently, exploiting the characteristics of modern hardware has become an important topic of database systems research. The aim of DaMoN 2010 is to bring together researchers who are interested in optimizing database performance on modern computing infrastructure by designing new data management techniques and tools.

Third International Workshop on Testing Database Systems (DBTest 2010)

http://www.cs.duke.edu/dbtest2010/
June 7, 2010
Location: Regency F

We are going through a period of rapid development in data management systems. On one side of the spectrum, commercial database system vendors are adding new features related to ease of management, semistructured data, data compression, parallelism, reliability, and security. On the other side, new data-processing systems are being developed over MapReduce backends, Key-Value stores, low-power devices, and widely-distributed systems.

New usage patterns, evolving hardware trends, and increased competition drive continuous innovation and expansion of these systems; both in terms of features as well as code size. As a result, it is increasingly expensive to test and tune these systems, and these stages tend to dominate the release cycle. It is not unusual to find fifty percent of the development cost being spent on testing and tuning, and several months being reserved for testing before a new release can be shipped. This situation will get worse in the future unless revolutionary new ideas are brought in.

DBTest 2010 continues the discussion between industry and academia in order to come up with a research agenda that describes important open problems in the area of testing database systems and applications. The long-term goal is to devise new techniques that solve these problems. These techniques will reduce the cost and time to test and tune database products so that users and vendors can spend more time and energy on actual innovations.

2010 DataBase MEntoring workshop (DB ME 2010)
http://www.cs.ubc.ca/~rap/dbme/
June 11, 2010
Location: Discovery

The goal of the workshop is to increase the diversity of the database community through mentoring women and underrepresented minorities. It is part of a series of *Discipline-specific Mentoring Workshops* organized by *Computing Research Association's* Committee on the *Status of Women in Computing* and the *Coalition to Diversify Computing (CDC).*

Ph.D. Workshop on Innovative Database Research (IDAR 2010)
http://www2.dcc.ufmg.br/eventos/idar2010/doku.php
June 11, 2010
Location: Vision

The Ph.D. Workshop on Innovative Database Research, held in conjunction with SIGMOD, brings together Ph.D. students working on topics related to the SIGMOD conference series. The workshop will offer Ph.D. students the opportunity to present, discuss, and receive feedback on their research in a constructive and international atmosphere. Prominent professors, researchers, and practitioners in the fields of database technology will participate actively in the discussions at the workshop.

PODS ORGANIZATION

PODS Executive Committee

Jan Paredaens (Chair, University of Antwerp, Belgium)
Georg Gottlob (University of Oxford, UK)
Phokion Kolaitis (UC Santa Cruz & IBM Research, USA)
Maurizio Lenzerini (Sapienza University of Rome, Italy)
Jianwen Su (UC Santa Barbara, USA)
Dirk Van Gucht (Indiana University, USA)

PODS General Chair

Jan Paredaens (University of Antwerp, Belgium)

PODS Program Chair

Dirk Van Gucht (Indiana University, USA)

PODS Proceedings Chair

George H. L. Fletcher (Eindhoven Univ. of Technology, The Netherlands)

PODS Publicity Chair

Yuqing Melanie Wu (Indiana University, USA)

PODS Program Committee

Pankaj K. Agarwal (Duke University, USA)
Marcelo Arenas (Pontifical Catholic Univ. of Chile)
Mikhail Atallah (Purdue University, USA)
Henrik Björklund (Umeå University, Sweden)
Mikołaj Bojańczyk (University of Warsaw, Poland)
Rada Chirkova (North Carolina State University, USA)
Graham Cormode (AT&T Labs Research, USA)
Jan Hidders (Delft Univ. of Technology, The Netherlands)
Benny Kimelfeld (IBM Almaden Research Center, USA)
Tova Milo (Tel Aviv University, Israel)
Jeffrey Naughton (University of Wisconsin, USA)
Lucian Popa (IBM Almaden Research Center, USA)
Riccardo Rosati (Sapienza University of Rome, Italy)
Dan Suciu (University of Washington, USA)
Val Tannen (University of Pennsylvania, USA)
Dirk Van Gucht (Indiana University, USA)
Jan Van den Bussche (Hasselt University, Belgium)
Stijn Vansummeren (Université Libre de Bruxelles, Belgium)
Limsoon Wong (National University of Singapore)

SIGMOD ORGANIZATION

SIGMOD General Chairs
Ahmed Elmagarmid (Purdue University)

SIGMOD Program Chair
Divyakant Agrawal (University of California at Santa Barbara)

SIGMOD Finance Chair
Marianne Winslett (University of Illinois at Urbana Champagne)

SIGMOD Sponsorship Chair
Michael Carey (University of California at Irvine)

SIGMOD Publicity Chairs
Sourav Bhowmick (NTU Singapore)
Kristen LeFevre (University of Michigan)

SIGMOD Proceedings Chair
Mohamed Mokbel (University of Minnesota)

SIGMOD Workshops Chair
Christian S. Jensen (Aalborg University)

SIGMOD Industrial Program Chair
Berthold Reinwald (IBM Almaden)

SIGMOD Tutorials Chair
Bernhard Mitschang (Universitat Stuttgart)

SIGMOD Keynotes & Panels Chair
Alon Halevy (Google, Inc.)

SIGMOD Demonstration Chair
Yufei Tao (Chinese University of Hong Kong)

SIGMOD Exhibits Chair
Luna Dong (AT&T Research)

SIGMOD New Researcher Symposium Chairs
Christoph Koch (Cornell University)
Yanlei Dao (University of Massachusetts)

Local Arrangements Chairs
Mourad Ouzzani (Purdue University)
Beth Plale (Indiana University)
Yuni Xia (Purdue University Indianapolis)
Local Workshop Chairs
Yuqing Melanie Wu (Indiana University)

Registration Chair
Hakan Ferhatosmanoglu (Ohio State University)

SIGMOD Undergraduate Research Program Chair
Sihem Amer-Yahia (Yahoo, Inc.)

SIGMOD New Initiatives Committee
Amr El Abbadi (University of California at Santa Barbara)
Jeffrey F. Naughton (University of Wisconsin at Madison)
Jennifer Widom (Stanford University)

SIGMOD Demonstrations Local Arrangements Chair

Brandeis H. Marshall (Purdue University)

SIGMOD Web/Information Chair

Wai Gen Yee (Illinois Institute of Technology)

SIGMOD Program Committee

Group Leaders

Gustavo Alonso (ETH Zurich)
Selcuk Candan (Arizona State University)
Graham Cormode (AT&T Research)
Susan Davidson (University of Pennsylvania)
Nick Koudas (University of Toronto)
Chen Li (University of California at Irvine)
David Lomet (Microsoft Research)
Sunil Prabhakar (Purdue University)
Evaggelia Pitoura (University of Ioannina)
Rajeev Rastogi, Yahoo Research India
Timos Sellis (IMIS/RC Athena and National Technical University of Athens)
Cyrus Shahabi (University of Southern California)

Committee Members

Daniel Abadi (Yale University)
Walid Aref (Purdue University)
Magdalena Balazinska (University of Washington)
Wolf-Tilo Balke (TU Braunschweig)
Nagender Bandi (Oracle, USA)
Abraham Bernstein (University of Zurich)
Phil Bernstein (Microsoft Research)
Subhash Bhalla (University of Aizu, Japan)
Peter Boncz (Centrum Wiskunde & Informatica)
Angela Bonifati (ICAR-CNR)
Fabio Casati (University of Trento)
Tiziana Catarci (University of Roma)
Stefano Ceri (Politecnico di Milano)
Kaushik Chakrabarti (Microsoft Research)
Chee Yong Chan (National University of Singapore)
Lei Chen (Hong Kong University of Science and Technology)
Ming-Syan Chen (National Taiwan University)
Yi Chen (Arizona State University)
Laura Chiticariu (IBM Almaden)
Panos Chrysanthis (University of Pittsburgh)
Chris Clifton (Purdue University)
Bin Cui (Peking University, China)
Sabrina De Capitani di Vimercati (Universita' degli Studi di Milano)
Amol Deshpande (University of Maryland)
Alin Dobra (University of Florida)
Christos Faloutsos (Carnegie-Melon University)
Wenfei Fan (University of Edinburgh)
Alan Fekete (University of Sydney)
Hakan Ferhatosmanoglu (Ohio State University)
Sumit Ganguly (IIT Kanpur, India)
Phil Gibbons (Intel Research)
Oliver Guenther (Humboldt University)
Sudipto Guha (University of Pennsylvania)
Dimitrios Gunopulos (University of Athens)
Hakan Hacigumus (NEC Research Laboratories of America)
Marios Hadjieleftherin (AT&T Shannon Labs)
Theo Haerder (University of Kaiserslautern)
Stavros Harizopoulos (HP Labs)
Wynne Hsu (National University of Singapore)
Christopher Jermaine (Rice University)

Theodore Johnson (AT&T Research)
Ben Kao (University of Hong Kong)
Daniel Keim (University of Kostanz)
Daniel Kifer (Pennsylvania State University)
Christoph Koch (Cornell University)
Yannis Kotidis (Athens University of Economics and Business)
Hans-Peter Kriegel (Ludwig-Maximilians-Universitaet Muenchen)
Laks V.S. Lakshmanan (University of British Columbia)
Jiaheng Lu (Renmin University of China)
Qiong Luo (Hong Kong University of Science & Technology)
Stefan Manegold (CWI: Centrum Wiskunde & Informatica)
Sharad Mehrotra (University of California at Irvine)
Ahmed Metwally (Google, Inc.)
Prasenjit Mitra (Pennsylvania State University)
Bongki Moon (University of Arizona)
Shinichi Morishita (University of Tokyo)
Kyriakos Mouratidis (Singapore Management University)
Suman Nath (Microsoft Research)
Wolfgang Nejdl (University of Hanover)
Thomas Neumann (Max-Planck Institute, Germany)
Raymond Ng (University of British Columbia)
Christopher Olston (Yahoo! Research)
Mourad Ouzzani (Purdue University)
Fatma Ozcan (IBM Almaden)
Hwee Hwa Pang (Singapore Management University)
Yannis Papakonstantinou (University of California at San Diego)
Jian Pei (Simon Fraser University)
Lin Qiao (IBM Alamden)
Mirek Riedwald (Northeastern University)
Uwe Rohm (University of Sydney)
Kenneth Salem (University of Waterloo)
Pierangela Samarati (University of Milan)
Maria Luisa Sapino (Universita di Torino)
Ralf Schenkel (Max-Planck Institute, Germany)
Mehul Shah (HP Laboratories)
Heng Tao Shen (University of Queensland, Australia)
Liuba Shrira (Brandeis University)
Alkis Simitsis (HP Labs)
Ambuj Singh (University of California at Santa Barbara)
Yannis Sismanis (IBM Almaden Tresearch)
Il-Yeol Song (Drexel University)
Nisheeth Srivastava (Bell Laboratories)
Torsten Suel (Polytechnic Institute of NYU)
Kian-Lee Tan (National University of Singapore)
Wang-Chiew Tan (University of California at Santa Cruz)
Nesime Tatbul (ETH Zurich)
Jun Tatemura (NEC Research Laboratories of America)
Martin Theobald (Max-Planck Institute, Germany)
Anthony Tung (National University of Singapore)
Gottfried Vossen (University of Muenster)
Haixun Wang (Microsoft Research Asia)
Min Wang (IBM T. J. Watson Research Center)
Wei Wang (University of North Carolina at Chapel Hill)
Dong Xin (Microsoft Research)
Xifeng Yan (University of California at Santa Barbara)
Haruo Y Yokota (Tokyo Institute of Technology)
Jeffrey Xu Yu (Chinese University of Hong Kong)
Carlo Zaniolo (University of California at Los Angeles)
Aidong Zhang (SUNY Buffalo)

SIGMOD Demonstrations Committee

Carl-Christian Kanne (Univ. of Mannheim)
Henrique Andrade IBM
Lisa Amini (IBM)
Michalis Petropoulos (Univ. of Buffalo)
Mohamed F. Mokbel (Univ. of Minnesota)
Shimin Chen (Intel)
Sudarshan S. Chawathe (Univ. of Maine)
Themis Palpanas (Univ. of Trento)
Vasilis Vassalos (Athens University of Economics and Business)
Vladislav Shkapenyuk (AT&T)
Wook-Shin Han (Kyungpook National University)
Xiaokui Xiao (Nanyang Technological University)

SIGMOD Industrial Committee

Ben Reed (Yahoo! Research)
Chet Murthy (IBM Research)
Dave Campbell (Microsoft)
James Hamilton (Amazon)
Jussi Myllymaki (Google)
Wen-Syan Li (SAP, China)

SIGMOD Tutorials Committee

Reynold Cheng (University of Hong Kong)
Bernhard Mitschang (Universitaet Stuttgart)
Felix Naumann (Hasso-Plattner-Institut an der Universitaet Potsdam)
Mario Nascimento (University of Alberta)

SOCC ORGANIZATION

SOCC General Chair
Joseph M. Hellerstein, U. C. Berkeley

SOCC Program Chairs
Surajit Chaudhuri, Microsoft Research
Mendel Rosenblum, Stanford University

SOCC Treasurer
Brian Cooper, Yahoo! Research

SOCC Publicity Chair
Aman Kansal, Microsoft Research

SOCC Steering Committee
Phil Bernstein, Microsoft Research
Ken Birman, Cornell University
Joseph M. Hellerstein, U. C. Berkeley
John Ousterhout, Stanford University
Raghu Ramakrishnan, Yahoo! Research
Doug Terry, Microsoft Research
John Wilkes, Google

SOCC Program Committee

Anastasia Ailamaki, EPFL
Brian Bershad, Google
Ken Birman, Cornell University
Michael Carey, UC Irvine
Felipe Cabrera, Amazon
Jeff Chase, Duke
Byung-gon Chun, Intel Labs Berkeley
Dilma M da Silva, IBM
David Dewitt, Microsoft
Shel Finkelstein, SAP
Armando Fox, UC Berkeley
Tal Garfinkel, Stanford
Alon Halevy, Google
James Hamilton, Amazon
Jeff Hammerbacher, Cloudera
Joe Hellerstein, UC Berkeley
Alfons Kemper, Technische Universität München
Donald Kossman, ETH
Orran Krieger, Vmware
Jim Larus, Microsoft
Jeffrey Naughton, University of Wisconsin, Madison
Christopher Olston, Yahoo
Hamid Pirahesh, IBM
Raghu Ramakrishnan, Yahoo!
Venugopalan Ramasubramanian, Microsoft Research
Hovav Shacham, UCSD
Donovan Schneider, Salesforce.com
Andy Warfield, University of British Columbia
Hakim Weatherspoon, Cornell
John Wilkes, Google

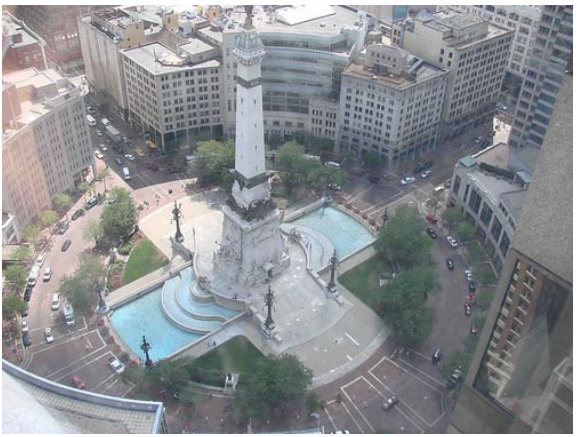

Indianapolis was selected as the site of the new state capital in 1820, the old state capitol having been Corydon since the formation of the state of Indiana. While most American state capitals tend to be located in the central region of their respective states, Indianapolis is the closest capital to being in the exact center of its state. It was founded on the White River both because of this, and because of the incorrect assumption that the river would serve as a major transportation artery. However, the waterway eventually proved to be too sandy for trade. Jeremiah Sullivan, a judge of the Indiana Supreme Court, invented the name Indianapolis by joining Indiana with polis, the Greek word for city; Indianapolis literally means "Indiana City". The state commissioned Alexander Ralston to design the new capital city. Ralston was an apprentice to the French architect Pierre L'Enfant, and he helped L'Enfant plan Washington, DC. Ralston's original plan for Indianapolis called for a city of only one square mile (3 km²). At the center of the city sat Governor's Circle, a large circular commons, which was to be the site of the governor's mansion. Meridian and Market Streets converge at the Circle and continue north-south and east-west, respectively. The Capitol moved from Corydon on January 10, 1825. The governor's mansion was eventually demolished in 1857 and in its place stands a 284-foot (87 m) tall neoclassical limestone and bronze monument, the Indiana Soldiers' and Sailors' Monument. The surrounding street is now known as Monument Circle.

High rise construction in Indianapolis started in 1888 with the Indiana Statehouse, and then the Soldiers' and Sailors' Monument in 1898. The 284-foot (87 m) tall monument sits at the center of Indianapolis and was the tallest structure in the city until the completion of City Hall in 1962. The Indiana Statehouse remained the second-tallest structure and the tallest building until 1962.

The city lies on the original east-west National Road. The first railroad to service Indianapolis, the Madison & Indianapolis, began operation on October 1, 1847, and subsequent railroad connections made expansive growth possible. Indianapolis was the home of the first Union Station, or common rail passenger terminal, in the United States. By the turn of the century, Indianapolis had become a large automobile manufacturer, rivaling the likes of Detroit. With roads leading out of the city in all directions, Indianapolis became a major hub of regional transport connecting to Chicago, Louisville, Cincinnati, Columbus, Detroit, Cleveland and St. Louis, befitting the capital of a state whose nickname is "The Crossroads of America."

Hyatt Restaurants

Eagle's Nest
Enjoy a romantic evening with exceptional cuisine at the only revolving restaurant in Indiana.

One South
Elevate the experience at Hyatt Regency Indianapolis' restaurant, One South.

Level One
Get together with friends at Hyatt Regency Indianapolis' popular lounge. Visit Hyatt's new lobby level lounge Level One.

Starbuck's® Coffee Shop
No matter what hour of the day, there is nothing like a fresh, hot cup of java to get you going or to rejuvenate the weary.

Restaurants, Brewery's and Clubs close to the Hyatt:
There are many outdoor dining areas around the hotel. The Hyatt consigner will assist you with cuisine and reservations.

Ameer Middle Eastern Cuisine
Indianapolis City Market
222 E. Market St., #33
Indianapolis (317) 681-8444

India Garden
209 N. Delaware St.
Indianapolis (317) 634-6060

Fogo de Chao
117 East Washington Street
Indianapolis (317) 638-4000

St Elmo Steak House
127 South Illinois Street
Indianapolis (317) 635-0636

Ruth's Chris Steak House
45 S. Illinois Street
Circle Centre Mall
Indianapolis (317) 633-1313

Oceanaire Seafood Room
30 South Meridian Street
Indianapolis (317) 955-2277

P F Chang China Bistro
49 West Maryland Street
Indianapolis (317) 974-5747

Palomino- Mediterranean
49 West Maryland Street
Indianapolis (317) 974-0400

Shapiro's Delicatessen
808 S Meridian St
Indianapolis (317) 631-4041

Greek Islands Restaurant
906 South Meridian Street
Indianapolis (317) 636-0700

Ram Restaurant big Horn Brewery
140 South Illinois Street
Indianapolis (317) 955-9900

Alcatraz Brewing Co.
49 West Maryland Street
Indianapolis (317) 488-1230

Rock bottom Brewery
10 W Washington St
Indianapolis (317) 681-8180

Rathskeller
401 East Michigan Street
Indianapolis (317) 636-0396

Nicky Blaine's
20 North Meridian Street
Indianapolis (317) 638-5588

Slippery Noodle
372 South Meridian Street
Indianapolis (317) 631-6968

Jazz Kitchen
5377 N College Ave, Indianapolis (317) 253-4900

Chatterbox Jazz Club
435 Massachusetts Ave, Indianapolis, IN 46204
(317) 636-0584

Howl at the Moon
American Cuisine with Dueling Grand Pianos
20 E. Georgia St., Indianapolis (317) 955-0300

Canal Café & Terrace at Indiana State Museum
650 W. Washington St., Indianapolis (317) 232-1637

Buca di Beppo
35 N. Illinois Street
Indianapolis (317) 632-2822

Events:

❖ June 8 - 11, 2010 - Pawtucket Red Sox vs. Indianapolis Indians at Victory Field
The Indianapolis Indians, Triple-A affiliate of the Pittsburgh Pirates, play home games at Victory Field, recognized as the "Best Minor League Ballpark in America." 501 W. Maryland St. Tickets $9-14. http://www.minorleaguebaseball.com/app/index.jsp?sid=t484

❖ All Summer - Butterflies at the Indianapolis Zoo
Fluttering, colorful and much-anticipated - the very popular Butterflies show returns to White River Gardens at the Indianapolis Zoo. Thousands of free-flying exotic and native butterflies, combined with the gentle sounds of flowing water and soothing music, create a calming ambiance inside the Hilbert Conservatory. It's a unique experience in beauty and relaxation, right in the heart of Indianapolis. 1200 W. Washington St. http://www.indyzoo.com/

❖ June 4, 2010 - IMA Summer Nights Film Series: Across the Universe at the Indianapolis Museum of Art. See movies every summer Friday on the IMA Amphitheater, located on the west side of the Museum. 4000 N. Michigan Rd. Tickets: $5, members; $10, public; free, children six and under. Films will be canceled in the event of inclement weather. http://www.imamuseum.org/blog/2010/04/01/summer-nights/

❖ June 5-6, 2010 - Indiana Festival at Conner Prairie
Discover Indiana's rich cultural heritage as you laugh, dance, stomp, stare, sway and weave your way through this fun-filled weekend. Immerse yourself in the diverse and wonderful music, art, dances and foods that make us who we are today. 13400 Allisonville Rd. Cost: $13, adults; $12, seniors 65+; $9, youth 2-12; free, members and youth under 2. http://connerprairie.org/

❖ June 5, 2010 - Brad Paisley at the Verizon Wireless Music Center
12880 E. 146th St. Tickets: $30.25, $35.75, $59.75. http://www.livenation.com/

❖ June 5, 2010 - The Tides at Indianapolis Artsgarden/Visitor Center
Indy's newest wacky, crazy, variety, Caribbean party band. FREE. 100 W. Washington St. http://www.thetidesband.com/

❖ June 5, 2010 - New York Liberty vs. Indiana Fever at Conseco Fieldhouse
The defending Eastern Conference champion Indiana Fever offers the excitement and skill of women's professional basketball in a family-friendly environment. Returning from a WNBA Finals appearance in 2009 and vying for its sixth consecutive playoff berth, the Fever enters the new season under the direction of third-year coach Lin Dunn. One Conseco Ct., 125 S. Pennsylvania St. http://feverbasketball.com/

❖ June 5, 2010 - Vintage Indiana Wine and Food Festival at Military Park
Visit the eleventh annual Vintage Indiana Wine & Food Festival, an award-winning festival promoting Indiana wine and food. Celebrate with the Indiana wineries as they offer samples of more than 100 award-winning wines (adult sampler ticket only—ID required—first 10,000 receive etched wine glass!) Enjoy delicious food from some of Indiana's top culinary experts. Visit the Wine & Food Experience Pavillion for food and wine sessions with chefs and winemakers. Tap your foot to some of Indiana's most popular musicians (Jennie DeVoe, Jon McLaughlin and Needtobreathe) LIVE throughout the day. Browse through the artists booths. Take your kids to the KidZone for some interactive fun. 601 W. New York St. http://www.vintageindiana.com/

❖ June 8, 2010 - Luke Austin Daugherty at the Indianapolis Artsgarden/Visitor Center
Acoustic singer and songwriter. FREE. 100 W. Washington St. http://visitindy.com/indianapolis/servlet/JspFramework?pub=/indianapolis/web/jsp/whattodo/detail.jsp&click=7017109:detail

❖ June 8, 2010 - International Violin Competition of Indianapolis Laureate Series: 1986 Gold Medalist Kyoko Takezawa in recital with Pianist Akir at the Indiana History Center 1986 Gold Medalist Kyoko Takezawa and pianist Akira Eguchi present a recital featuring Mendelssohn's Violin Sonata in F major, Brahms' Violin Sonata No. 2 in A major, and Poulenc's Violin Sonata. As the embodiment of musicality, violinist Kyoko Takezawa electrifies audiences with a richness of playing, a virtuosic confidence of feeling, and a fiery intensity that establishes her as one of today's foremost violinists. Ms. Takezawa's interpretive insight and indisputable talent have made her a sought-after soloist with many of the world's leading orchestras. Acclaimed for his extraordinary artistry, maturity and intelligence, Akira Eguchi has captivated audiences and critics throughout the world as a piano soloist, chamber musician, harpsichord player and collaborative pianist. He was praised as a "pianist of fluency and rectitude" by The New York Times. 450 W. Ohio St. Tickets: $22, adults; $17, seniors age 60 and above; $10, students. http://www.violin.org/

❖ June 9 - July 3, 2010 - Jersey Boys at The Murat Centre "Too good to be true!" raves the New York Post for Jersey Boys, the 2006 Tony Award®-winning Best Musical about Rock and Roll Hall of Famers The Four Seasons: Frankie Valli, Bob Gaudio, Tommy DeVito and Nick Massi. This is the story of how four blue-collar kids became one of the greatest successes in pop music history. They wrote their own songs, invented their own sounds and sold 175 million records worldwide – all before they were 30! Jersey Boys, winner of the 2006 Grammy® Award for Best Musical Show Album, features their hit songs "Sherry," "Big Girls Don't Cry," "Rag Doll," "Oh What a Night" and "Can't Take My Eyes Off You." "It will run for centuries!" proclaims Time Magazine. The Jersey Boys creative team comprises two-time Tony Award®-winning director Des McAnuff, book writers Marshall Brickman and Rick Elice, composer Bob Gaudio, lyricist Bob Crewe and choreography by Sergio Trujillo. 502 N. New Jersey St. http://www.ticketmaster.com/venueartist/41109/938480

❖ June 10 - July 18, 2010 - High School Musical at Beef & Boards Dinner Theatre
It's the Disney sensation for a whole new generation: High School Musical comes to life on stage! Visit East High School where Troy and Gabriella reunite after a chance meeting on their winter break. Although they have different interests, they share a talent for singing and decide to audition for their school musical. It almost costs them their friendships, but brings everyone together in the end. 9301 Michigan Rd. Tickets: $35-$58. Includes a dinner buffet, full fruit and salad bar and unlimited coffee, tea and lemonade. http://www.beefandboards.com/

❖ June 10-13, 2010 - Czech Collectors Association at the Hilton Indianapolis North Hotel
The Czech Collectors Association's annual convention will include the following: Tom and Jane Rood will be presenting their Czech ladies to us. Mary Gawle will discuss Documenting and Insuring Your Collection. Jiří Harcuba, the best glass engraver in the world, will be speaking and we are hoping to arrange some equipment so he can do demonstrations. 8181 N. Shadeland Ave. http://www.czechcollectors.org/

❖ June 10, 2010 - Melissa Gallant at the Indianapolis Artsgarden/Visitor Center
Harpist performing a variety of classical and contemporary selections. 100 W. Washington St. http://www.melissagallant.com/

❖ June 11, 12 & 13, 2010 - Goodguys Hot Rod Nationals at the O'Reilly Raceway Park at Indianapolis. The Goodguys Hot Rod Nationals returns for the 22nd yearly summer visit featuring: The largest nostalgia drag race and show featuring over 4,000 hot rods, customs and classics, Hundreds of vendors and commercial exhibitors on the manufacturer midway, Three big days of nostalgia drag racing action and show n' shine. 10267 U.S. Hwy 136. Cost: $20, adults; $6, kids 7-12; free, kids 6 and under. http://www.good-guys.com/

❖ June 11, 12 & 13, 2010 - Indianapolis Symphony Orchestra Pops Series: Generations of Rock with Michael Cavanaugh Jack Everly, Conductor. Michael Cavanaugh, Vocals and Piano.

Back by popular demand, Michael Cavanaugh and his band return to the Hilbert Circle Theatre to perform rock hits from the '50s, '60s and '70s. If you like the music of Jerry Lee Lewis, the Eagles, Led Zeppelin, Elton John, Billy Joel and Bryan Adams, then you'll love this fresh and dynamic interpretation of their greatest hits. 45 Monument Circle. http://www.indianapolissymphony.org/

❖ June 11 & 12, 2010 - Reincarnation at the Pike Performing Arts Center
Award winning choreographer Gregory Hancock has been creating extraordinary dance works for well over 15 years. With over 100 repertoire pieces, Mr. Hancock continues to be prolific each and every season. With new pieces being constantly developed, Gregory Hancock Dance Theatre rarely has the time to revisit the older, classic repertoire. This June, GHDT is very excited to look back and revive some of the greatest works from their distant past. Vintage made Vanguard! 6701 Zionsville Rd. http://www.gregoryhancockdancetheatre.org/

❖ June 11, 2010 - IMA Summer Nights Film Series: Rebel Without a Cause at the Indianapolis Museum of Art See movies every summer Friday on the IMA Amphitheater, located on the west side of the Museum. 4000 N. Michigan Rd. Tickets: $5, members; $10, public; free, children six and under. Films will be canceled in the event of inclement weather. http://www.imamuseum.org/blog/2010/04/01/summer-nights/

❖ June 11, 2010 - Zoobilation 2010 at the Indianapolis Zoo
Delicious gourmet goodies from over 70 area restaurants highlight Zoobilation, along with top notch entertainment on stages that stretch from the Washington Street Bridge to all parts of Zoo grounds. Each year, Zoobilation raises more than a million dollars to support the care and feeding of the Indianapolis Zoo's living collection of 2,500 animals and 6,000 plants. Zoobilation proceeds also support the Zoo's internationally significant conservation and research projects. 1200 W. Washington St. http://www.indyzoo.com/

Museums/Gardens:

- Eiteljorg Museum, 500 W. Washington St. http://www.eiteljorg.org/
- Indiana State Museum/IMAX Theatre, 650 W. Washington St. http://www.indianamuseum.org/
- NCAA Hall of Champions, 700 W. Washington St. http://www.ncaahallofchampions.org/
- Indianapolis Zoo, 1200 W. Washington St. http://www.indyzoo.com/
- White River State Park, 801 W. Washington St. http://www.inwhiteriver.org/
- The Children's Museum of Indianapolis, 3000 N. Meridian St. http://www.childrensmuseum.org/
- Indianapolis Museum of Art, 4000 N. Michigan Rd. http://www.imamuseum.org/
- Soldiers & Sailors Monument/Monument Circle, 1 Monument Circle. http://www.in.gov/iwm/civilwar/
- Indianapolis Artsgardent/Visitor Center, 100 W. Washington St. http://www.indyarts.org/
- Indiana War Memorial Plaza Historic District, 400 N. Meridian St. http://visitindy.com/indianapolis/web/jsp/whattodo/results.jsp?c=attractionsmuseums:all&m=v&p=1